VISIT…

Chapter 11

Materials processing and design

11.1 Introduction and synopsis

A *process* is a method of shaping, finishing or joining a material. *Sand casting, injection moulding, polishing* and fusion *welding* are all processes; there are hundreds of them. It is important to choose the right process-route at an early stage in the design before the cost-penalty of making changes becomes large. The choice, for a given component, depends on the material of which it is to be made, on its size, shape and precision, and on how many are to be made — in short, on the *design requirements.* A change in design requirements may demand a change in process route.

Each process is characterized by a set of *attributes:* the materials it can handle, the shapes it can make and their precision, complexity and size. The intimate details of processes make tedious reading, but have to be faced: we describe them briefly in the following section, using Process Selection Charts to capture their attributes. *Process selection* is the act of finding the best match between process attributes and design requirements.

Methods for doing this are developed in the remaining sections of this chapter. In using them, one should not forget that material, shape and processing interact (Figure 11.1). Material properties and shape limit the choice of process: ductile materials can be forged, rolled and drawn; those which are brittle must be shaped in other ways. Materials which melt at modest temperatures to low-viscosity liquids can be cast; those which do not have to be processed by other routes. Slender shapes can be made easily by rolling or drawing but not by casting. High precision is possible by machining but not by forging, and so on. And processing affects properties. Rolling and forging change the texture of metals and align the inclusions they contain, enhancing strength and ductility. Composites acquire their properties during processing by control of lay-up; for these the interactions between function, material, shape and process are particularly strong.

Like the other aspects of design, process selection is an iterative procedure. The first iteration gives one or more possible processes-routes. The design must then be re-thought to adapt it, as far as possible, to ease of manufacture by the most promising route. The final choice is based on a comparison of *process cost,* requiring the use of cost models developed later in this chapter, and on *supporting information:* case histories, documented experience and examples of process-routes used for related products.

11.2 Processes and their influence on design

Now for the inevitable catalogue of manufacturing processes. It will be kept as concise as possible; details can be found in the numerous books listed in Further reading at the end of this chapter.

Manufacturing processes can be classified under the nine headings shown in Figure 11.2. *Primary processes* create shapes. The first row lists five primary forming processes: casting, moulding,

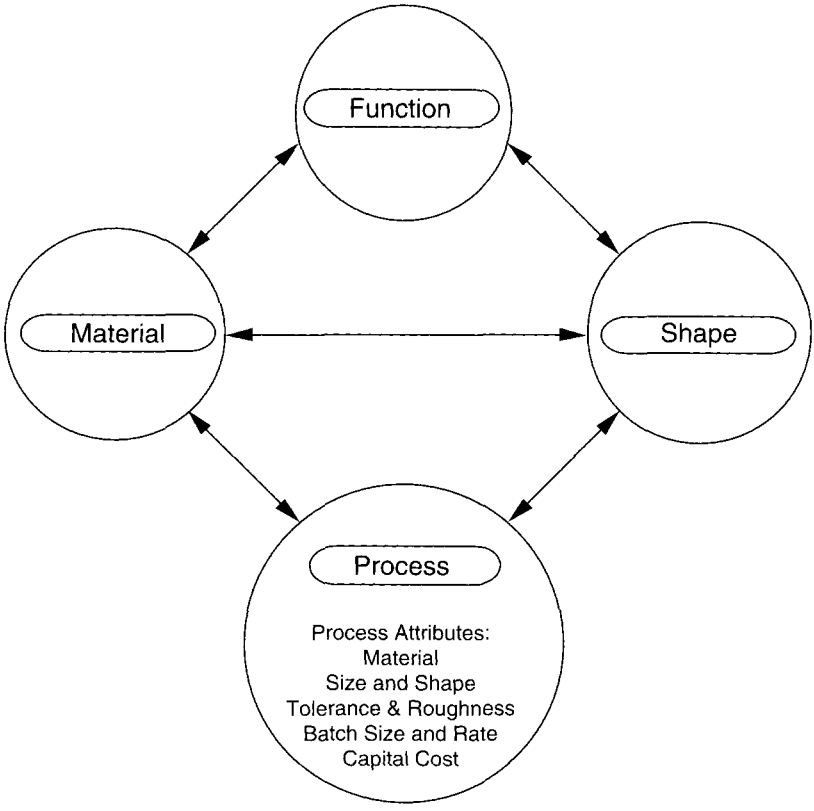

Fig. 11.1 Processing selection depends on material and shape. The 'process attributes' are used as criteria for selection.

deformation, powder methods, methods for forming composites, special methods and rapid prototyping. *Secondary processes* modify shapes; here they are shown collectively as 'machining'; they add features to an already shaped body. These are followed by *tertiary processes*: like heat treatment, which enhance surface or bulk properties. The classification is completed by *finishing* and *joining*.

(a) In *casting*, a liquid is poured into a mould where it solidifies by cooling (metals) or by reaction (thermosets). Casting is distinguished from moulding, which comes next, by the low viscosity of the liquid: it fills the mould by flow under its own weight (gravity casting, Figure 11.3) or under a modest pressure (centrifugal casting and pressure die casting, Figure 11.4). Sand moulds for one-off castings are cheap; metal dies for making large batches can be expensive. Between these extremes lie a number of other casting methods: shell, investment, plaster-mould and so forth.

Cast shapes must be designed for easy flow of liquid to all parts of the mould, and for progressive solidification which does not trap pockets of liquid in a solid shell, giving shrinkage cavities. Whenever possible, section thicknesses are made uniform (the thickness of adjoining sections should not differ by more than a factor of 2). The shape is designed so that the pattern and the finished casting can be removed from the mould. Keyed-in shapes are avoided because they lead to 'hot tearing' (a tensile creep-fracture) as the solid cools and shrinks. The tolerance and surface finish

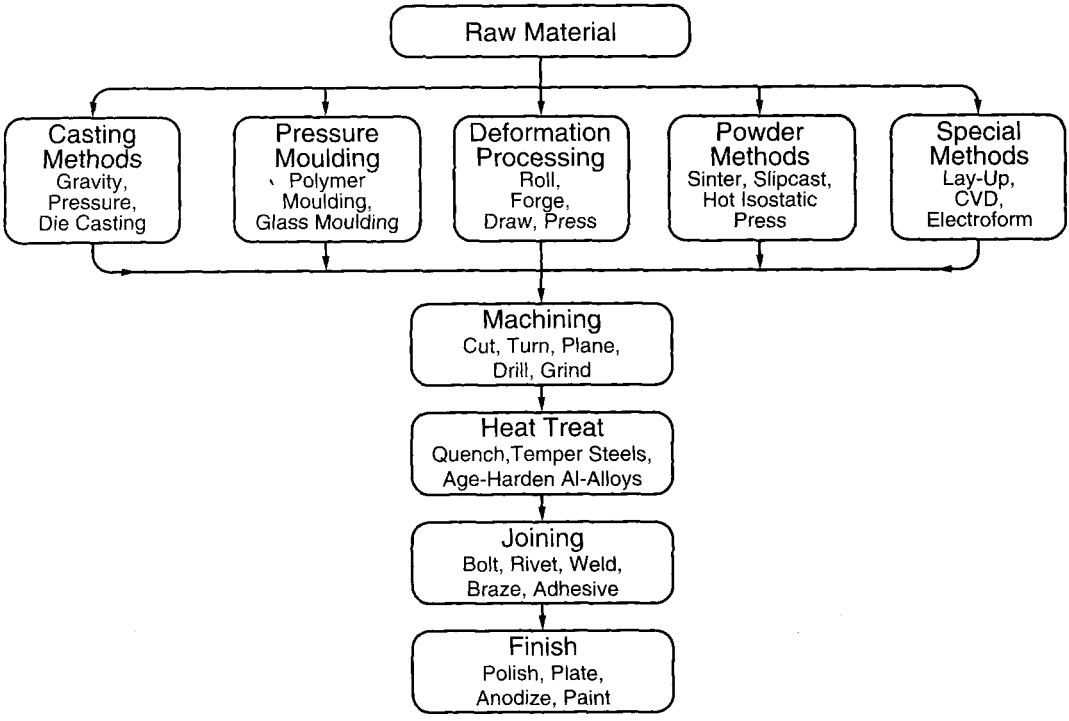

Fig. 11.2 The *nine* classes of process. The first row contains the primary shaping processes; below lie the secondary shaping, joining and finishing processes.

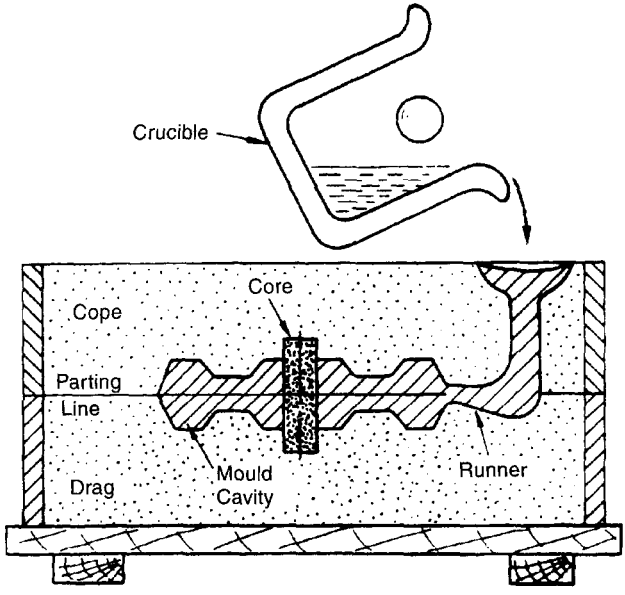

Fig. 11.3 Sand casting. Liquid metal is poured into a split sand mould.

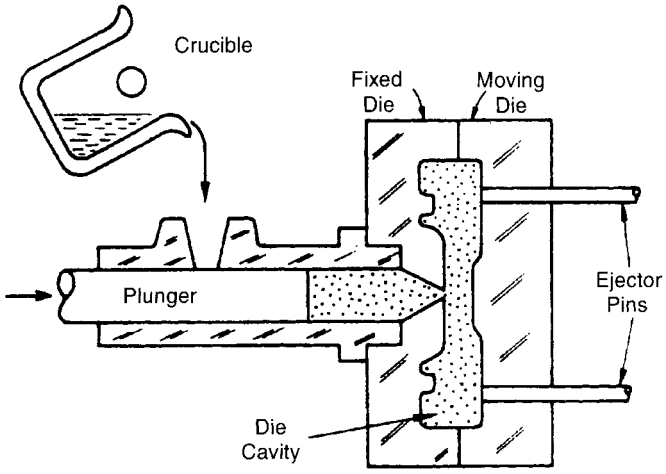

Fig. 11.4 Die casting. Liquid is forced under pressure into a split metal mould.

of a casting vary from poor for cheap sand-casting to excellent for precision die-castings; they are quantified at page 272.

(b) *Moulding* is casting, adapted to materials which are very viscous when molten, particularly thermoplastics and glasses. The hot, viscous fluid is pressed (Figure 11.5) or injected (Figures 11.6 and 11.7) into a die under considerable pressure, where it cools and solidifies. The die must withstand repeated application of pressure, temperature, and the wear involved in separating and removing the part, and therefore is expensive. Elaborate shapes can be moulded, but at the penalty of complexity in die shape and in the way it separates to allow removal.

Blow-moulding (Figure 11.8) uses a gas pressure to expand a polymer or glass blank into a split outer-die. It is a rapid, low-cost process well suited for mass-production of cheap parts like milk bottles.

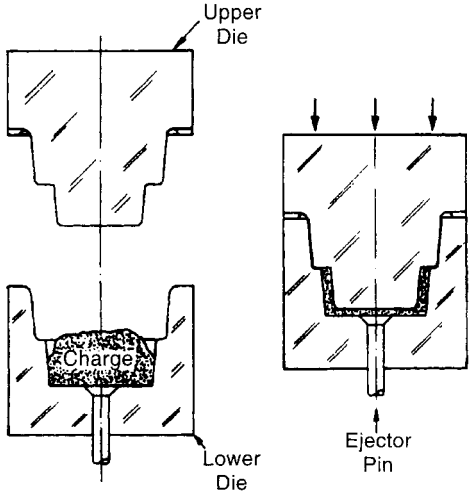

Fig. 11.5 Moulding. A hot slug of polymer or glass is pressed to shape between two dies.

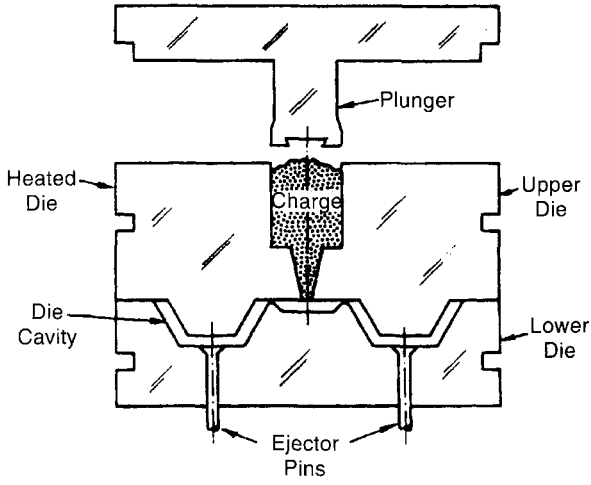

Fig. 11.6 Transfer-moulding. A slug of polymer or glass in a heated mould is forced into the mould cavity by a plunger.

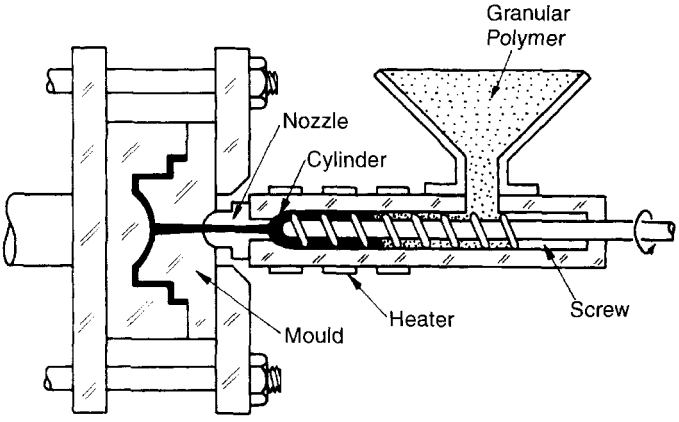

Fig. 11.7 Injection-moulding. A granular polymer (or filled polymer) is heated, compressed and sheared by a screw feeder, forcing it into the mould cavity.

(c) *Deformation processing* (Figures 11.9 to 11.12) can be hot, warm or cold. Extrusion, hot forging and hot rolling ($T > 0.55T_m$) have much in common with moulding, though the material is a true solid not a viscous liquid. The high temperature lowers the yield strength and allows simultaneous recrystallization, both of which lower the forming pressures. Warm working ($0.35T_m < T < 0.55T_m$) allows recovery but not recrystallization. Cold forging, rolling and drawing ($T < 0.35T_m$) exploit work hardening to increase the strength of the final product, but at the penalty of higher forming pressures.

Forged parts are designed to avoid rapid changes in thickness and sharp radii of curvature. Both require large local strains which can cause the material to tear or to fold back on itself ('lapping'). Hot forging of metals allows bigger changes of shape but generally gives a poor surface and

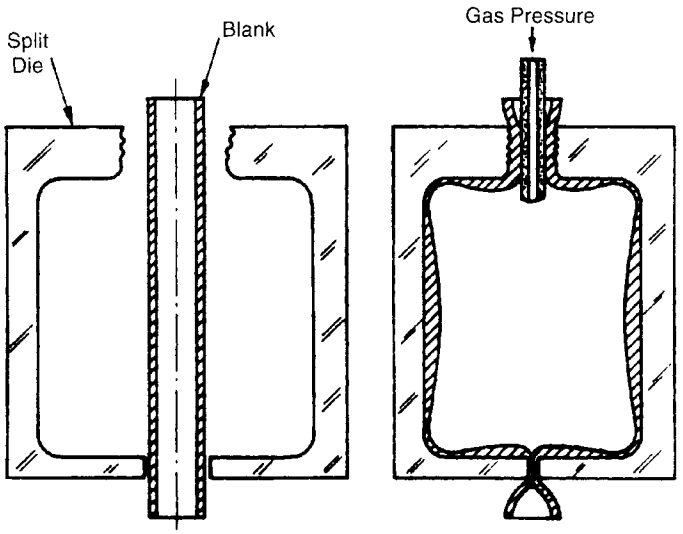

Fig. 11.8 Blow-moulding. A tubular or globular blank of hot polymer or glass is expanded by gas pressure against the inner wall of a split die.

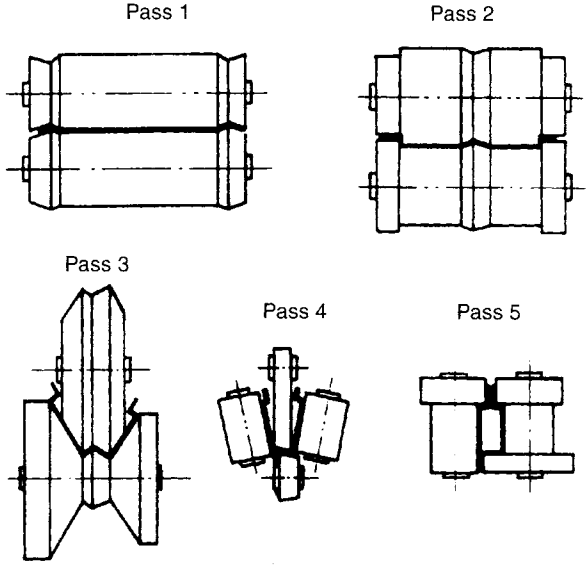

Fig. 11.9 Rolling. A billet or bar is reduced in section by compressive deformation between the rolls. The process can be hot ($T > 0.55T_m$), warm ($0.35T_m < T < 0.55T_m$) or cold ($T < 0.35T_m$).

tolerance because of oxidation and warpage. Cold forging gives greater precision and finish, but forging pressures are higher and the deformations are limited by work hardening.

Sheet metal forming (Figure 11.12) involves punching, bending, and stretching. Holes cannot be punched to a diameter less than the thickness of the sheet. The minimum radius to which a sheet can be bent, its *formability*, is sometimes expressed in multiples of the sheet thickness t: a value

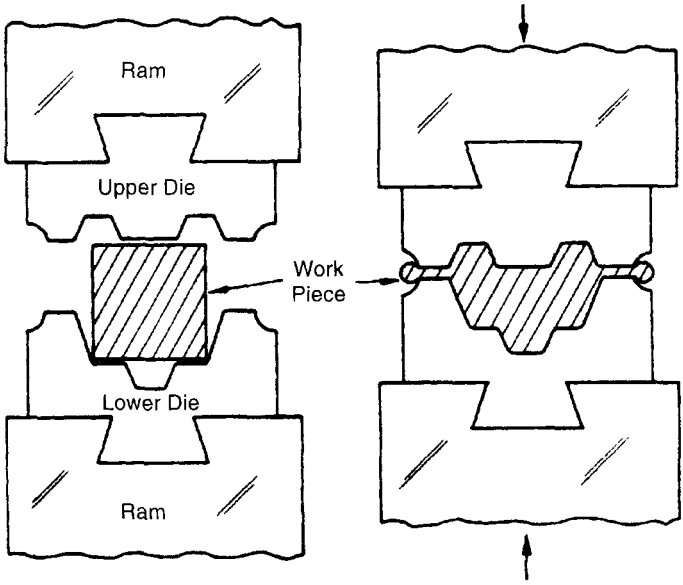

Fig. 11.10 Forging. A billet or blank is deformed to shape between hardened dies. Like rolling, the process can be hot, warm or cold.

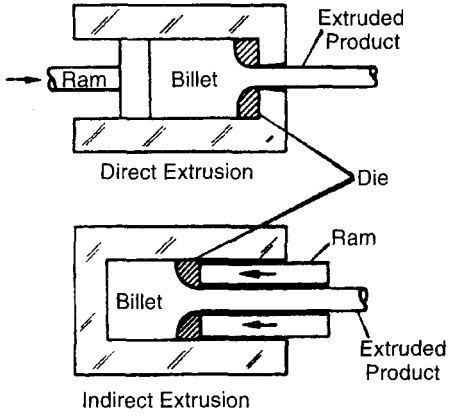

Fig. 11.11 Extrusion. Material is forced to flow through a die aperture to give a continuous prismatic shape. Hot extrusion is carried out at temperatures up to $0.9T_m$; cold extrusion is at room temperature.

of 1 is good; one of 4 is average. Radii are best made as large as possible, and never less than t. The formability also determines the amount the sheet can be stretched or drawn without necking and failing. The *limit forming diagram* gives more precise information: it shows the combination of principal strains in the plane of the sheet which will cause failure. The part is designed so that the strains do not exceed this limit.

(d) *Powder methods* create the shape by pressing and then sintering fine particles of the material. The powder can be cold-pressed and then sintered (heated at up to $0.8T_m$ to give bonding); it can

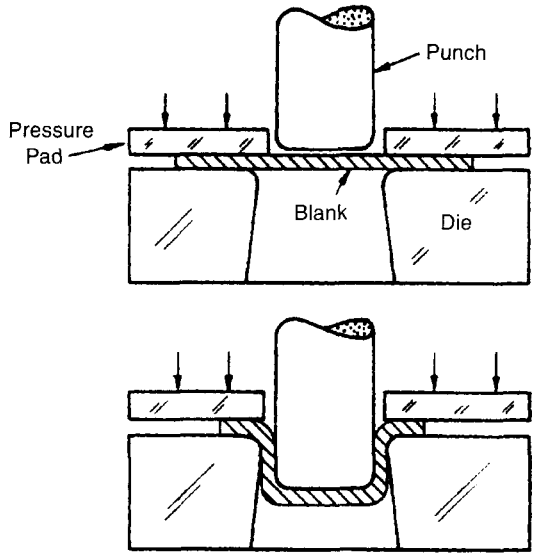

Fig. 11.12 Drawing. A blank, clamped at its edges, is stretched to shape by a punch.

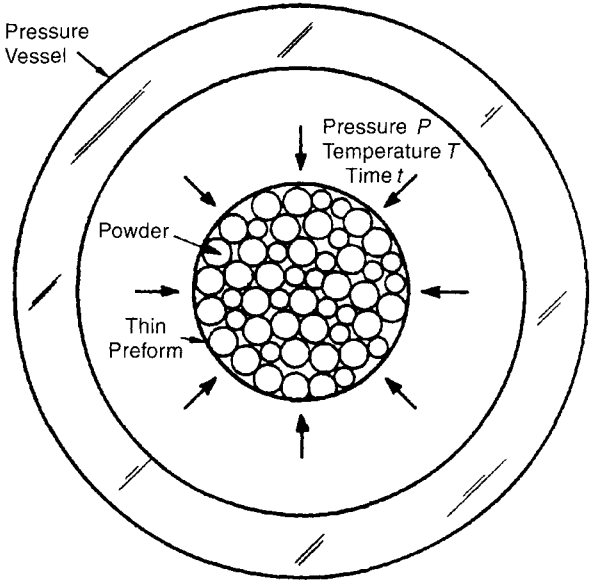

Fig. 11.13 Hot isostatic pressing. A powder in a thin, shaped, shell or preform is heated and compressed by an external gas pressure.

be pressed in a heated die ('die pressing'); or, contained in a thin preform, it can be heated under a hydrostatic pressure ('hot isostatic pressing' or 'HIPing', Figure 11.13). Metals and ceramics which are too high-melting to cast and too strong to deform can be made (by chemical methods) into powders and then shaped in this way. But the processes is not limited to 'difficult' materials; almost any material can be shaped by subjecting it, as a powder, to pressure and heat.

Powder pressing is most widely used for small metallic parts like gears and bearings for cars and appliances, and for fabricating almost all engineering ceramics. It is economic in its use of material, it allows parts to be fabricated from materials that cannot be cast, deformed or machined, and it can give a product which requires little or no finishing.

Since pressure is not transmitted uniformly through a bed of powder, the length of a die-pressed powder part should not exceed 2.5 times its diameter. Sections must be near-uniform because the powder will not flow easily round corners. And the shape must be simple and easily extracted from the die.

(e) *Composite fabrication* methods are adapted to make polymer-matrix composites reinforced with continuous or chopped fibres. Large components are fabricated by filament winding (Figure 11.14) or by laying-up pre-impregnated mats of carbon, glass or Kevlar fibre ('pre-preg') to the required thickness, pressing and curing. Parts of the process can be automated, but it remains a slow manufacturing route; and, if the component is a critical one, extensive ultrasonic testing may be necessary to confirm its integrity. So lay-up methods are best suited to a small number of high-performance, tailor-made, components. More routine components (car bumpers, tennis racquets) are made from chopped-fibre composites by pressing and heating a 'dough' of resin containing the fibres, known as bulk moulding compound (BMC) or sheet moulding compound (SMC), in a mould, or by injection moulding a rather more fluid mixture into a die as in Figures 11.5, 11.6 and 11.7. The flow pattern is critical in aligning the fibres, so that the designer must work closely with the manufacturer to exploit the composite properties fully.

(f) *Special methods* include techniques which allow a shape to be built up atom-by-atom, as in electro-forming and chemical and physical vapour deposition. They include, too, various spray-forming techniques (Figure 11.15) in which the material, melted by direct heating or by injection into a plasma, is sprayed onto a former — processes which lend themselves to the low-number production of small parts, made from difficult materials.

(g) *Machining* almost all engineering components, whether made of metal, polymer or ceramic, are subjected to some kind of machining (Figure 11.16) or grinding (a sort of micro-machining, as in Figure 11.17) during manufacture. To make this possible they should be designed to make gripping and jigging easy, and to keep the symmetry high: symmetric shapes need fewer operations. Metals differ greatly in their *machinability*, a measure of the ease of chip formation, the ability to give a smooth surface, and the ability to give economical tool life (evaluated in a standard test). Poor machinability means higher cost.

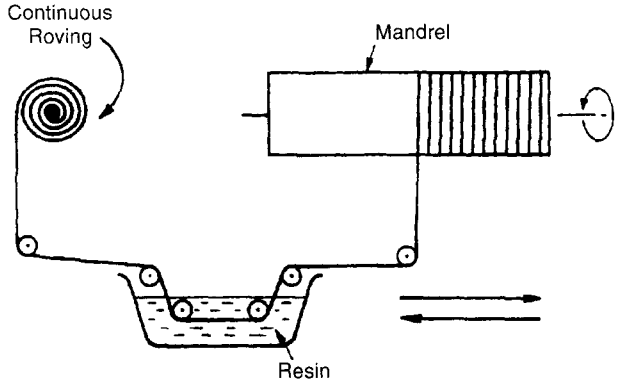

Fig. 11.14 Filament winding. Fibres of glass, Kevlar or carbon are wound onto a former and impregnated with a resin–hardener mix.

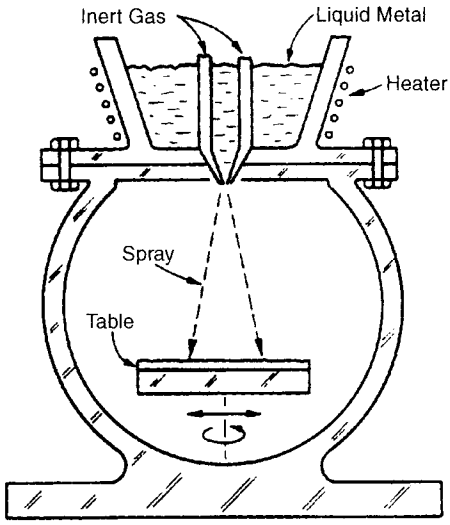

Fig. 11.15 Spray forming. Liquid metal is 'atomized' to droplets by a high velocity gas stream and projected onto a former where it splats and solidifies.

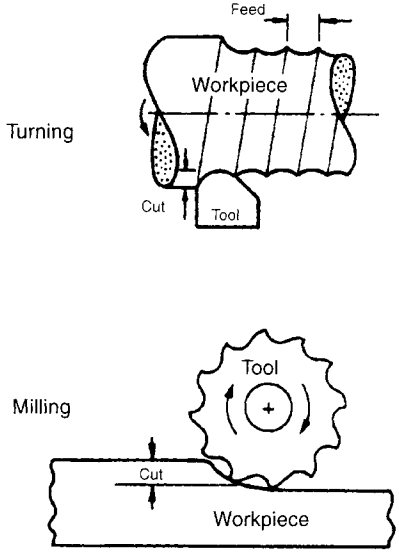

Fig. 11.16 Machining: turning (above left) and milling (below). The sharp, hardened tip of a tool cuts a chip from the workpiece surface.

Most polymers machine easily and can be polished to a high finish. But their low moduli mean that they deflect elastically during the machining operation, limiting the tolerance. Ceramics and glasses can be ground and lapped to high tolerance and finish (think of the mirrors of telescopes). There are many 'special' machining techniques with particular applications; they include electro-machining, spark machining, ultrasonic cutting, chemical milling, cutting by water-jets, sand-jets, electron beams and laser beams.

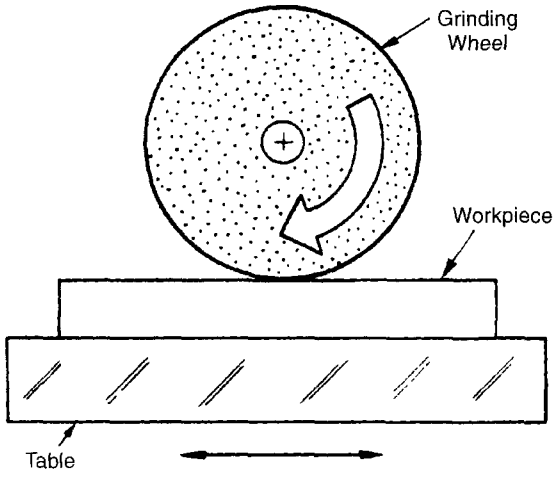

Fig. 11.17 Grinding. The cutting 'tool' is the sharp facet of an abrasive grain; the process is a sort of micro-machining.

Machining operations are often finishing operations, and thus determine finish and tolerance (pp. 271–2). Higher finish and tolerance mean higher cost; overspecifying either is a mistake.

(h) *Heat treatment* is a necessary part of the processing of many materials. Age-hardening alloys of aluminium, titanium and nickel derive their strength from a precipitate produced by a controlled heat treatment: quenching from a high temperature followed by ageing at a lower one. The hardness and toughness of steels is controlled in a similar way: by quenching from the 'austenitizing' temperature (about 800°C) and tempering.

Quenching is a savage procedure; thermal contraction can produce stresses large enough to distort or crack the component. The stresses are caused by a non-uniform temperature distribution, and this, in turn, is related to the geometry of the component. To avoid damaging stresses, the section should be as uniform as possible, and nowhere so large that the quench-rate falls below the critical value required for successful heat treatment. Stress concentrations should be avoided: they are usually the source of quench cracks. Materials which have been moulded or deformed may contain internal stresses which can be removed, at least partially, by stress-relief anneals — another sort of heat treatment.

(i) *Joining* is made possible by a number of techniques. Bolting and riveting (Figure 11.18), welding, brazing and soldering (Figure 11.19) are commonly used for metals. Polymers are joined by snap-fasteners (Figure 11.18 again), and by thermal bonding. Ceramics can be diffusion-bonded to themselves, to glasses and to metals. Improved adhesives give new ways of bonding all classes of materials (Figure 11.20). Friction welding (Figure 11.21) and friction-stir welding rely on the heat and deformation generated by friction to create a bond.

If components are to be welded, the material of which they are made must be characterized by a high *weldability*. Like *machinability*, it measures a combination of basic properties. A low thermal conductivity allows welding with a low rate of heat input, and gives a less rapid quench when the weld torch is removed. Low thermal expansion gives small thermal strains with less risk of distortion. A solid solution is better than an age-hardened alloy because, in the heat-affected zone on either side of the weld, overageing and softening can occur.

Welding always leaves internal stresses which are roughly equal to the yield strength. They can be relaxed by heat treatment but this is expensive, so it is better to minimize their effect by good

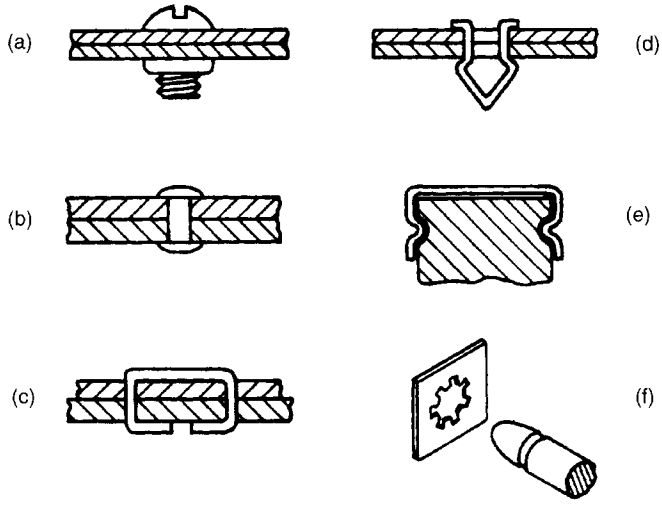

Fig. 11.18 Fasteners: (a) bolting; (b) riveting; (c) stapling; (d) push-through snap fastener; (e) push-on snap fastener; (f) rod-to-sheet snap fastener.

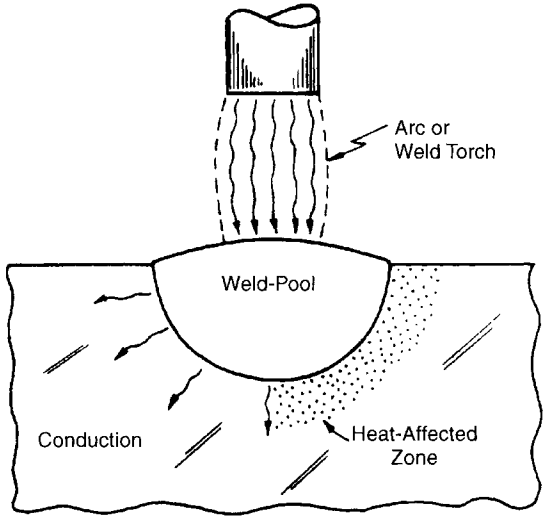

Fig. 11.19 Welding. A torch melts both the workpiece and added weld-metal to give a bond which is like a small casting.

design. To achieve this, parts to be welded are made of equal thickness whenever possible, the welds are located where stress or deflection is least critical, and the total number of welds is minimized.

The large-volume use of fasteners is costly because it is difficult to automate; welding, crimping or the use of adhesives can be more economical.

(j) *Finishing* describes treatments applied to the surface of the component or assembly. They include polishing, plating, anodizing and painting, they aim to improve surface smoothness, protect against corrosion, oxidation and wear, and to enhance appearance.

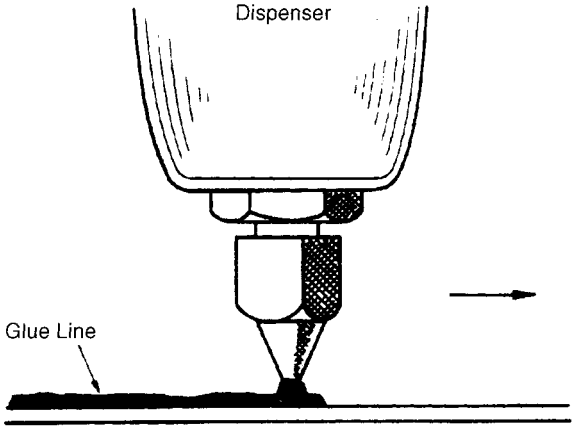

Fig. 11.20 Adhesive bonding. The dispenser, which can be automated, applies a glue-line onto the workpiece against which the mating face is pressed.

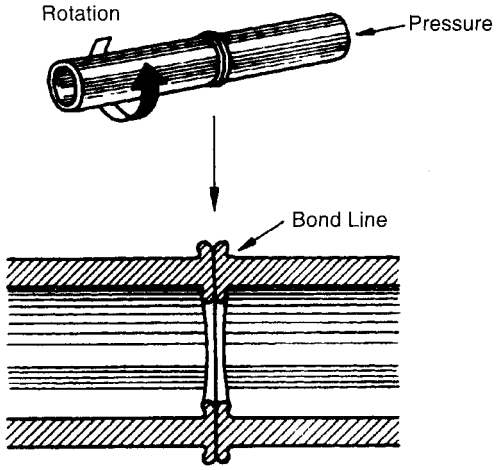

Fig. 11.21 Friction welding. A part, rotating at high speed, is pressed against a mating part which is clamped and stationary. Friction generates sufficient heat to create a bond.

Plating and painting are both made easier by a simple part shape with largely convex surfaces. Channels, crevices and slots are difficult to reach with paint equipment and often inadequately coated by electroplates.

(k) *Rapid prototyping systems* (RPS) allow single examples of complex shapes to be made from numerical data generated by CAD solid-modelling software. The motive may be that of visualization: the aesthetics of an object may be evident only when viewed as a prototype. It may be that of pattern-making: the prototype becomes the master from which moulds for conventional processing, such as casting, can be made. Or — in complex assemblies — it may be that of validating intricate geometry, ensuring that parts fit, can be assembled, and are accessible. All RPS can create shapes of great complexity with internal cavities, overhangs and transverse features, although the precision, at present, is limited to ± 0.3 mm at best.

The methods build shapes layer-by-layer, rather like three-dimensional printing, and are slow (typically 4–40 hours per unit). There are four broad classes of RPS.

(i) The shape is built up from a thermoplastic fed to a single scanning head which extrudes it like a thin layer of toothpaste ('Fused Deposition Modelling' or FDM), exudes it as tiny droplets ('Ballistic Particle Manufacture', BPM, Figure 11.22), or ejects it in a patterned array like a bubble-jet printer ('3-D printing').

(ii) Screen-based technology like that used to produce microcircuits ('Solid Ground Curing' or SGC, Figure 11.23). A succession of screens admits UV light to polymerize a photo-sensitive monomer, building shapes layer-by-layer.

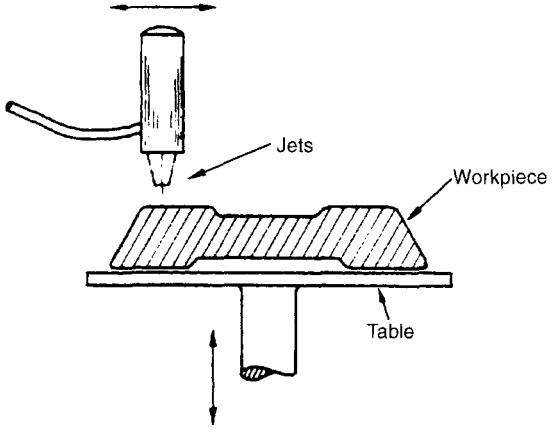

Fig. 11.22 Ballistic particle manufacture (BPM), a rapid prototyping method by which a solid body is created by layer-by-layer deposition of polymer droplets.

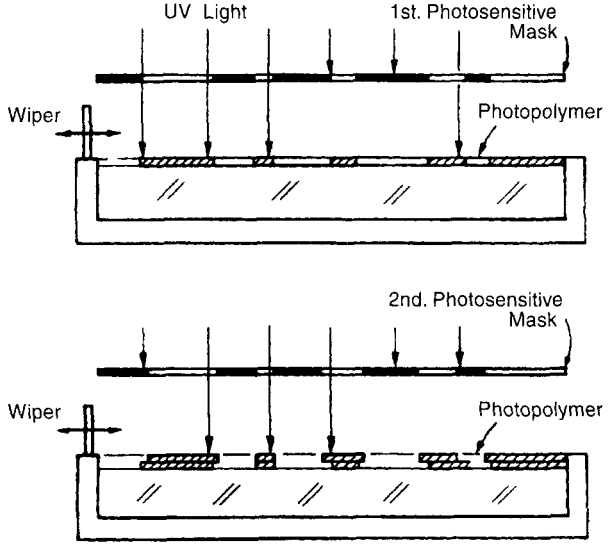

Fig. 11.23 Solid ground curing (SGC), a rapid prototyping method by which solid shapes are created by sequential exposure of a resin to UV light through glass masks.

(iii) Scanned-laser induced polymerization of a photo-sensitive monomer ('Stereo-lithography' or SLA, Figure 11.24). After each scan, the workpiece is incrementally lowered, allowing fresh monomer to cover the surface. Selected laser sintering (SLS) uses similar laser-based technology to sinter polymeric powders to give a final product. Systems which extend this to the sintering of metals are under development.

(iv) Scanned laser cutting of bondable paper elements (Figure 11.25). Each paper-thin layer is cut by a laser beam and heat bonded to the one below.

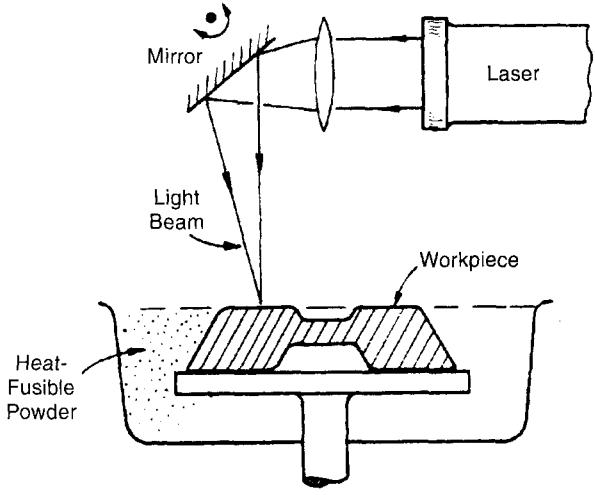

Fig. 11.24 Stereo-lithography (SLA), a rapid prototyping method by which solid shapes are created by laser-induced polymerization of a resin.

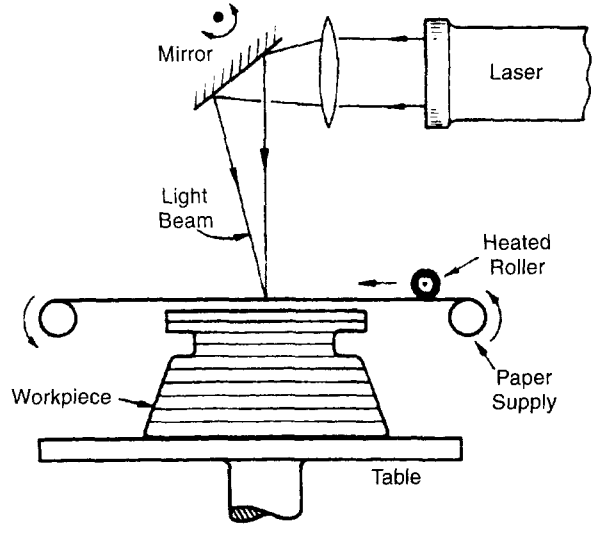

Fig. 11.25 Laminated object manufacture (LOM), a rapid prototyping method by which a solid body is created from layers of paper, cut by a scanning laser beam and bonded with a heat-sensitive polymer.

To be useful, the prototypes made by RPS are used as masters for silicone moulding, allowing a number of replicas to be cast using high-temperature resins or metals.

Enough of the processes themselves; for more detail the reader will have to consult the Further reading section.

11.3 Process attributes

The *kingdom* of processes can be classified in the way shown in top half of Figure 11.26. There are the broad *families*: casting, deformation, moulding, machining, compaction of powders, and such like. Each family contains many *classes*: casting contains sand-casting, die-casting, and investment casting, for instance. These in turn have many *members*: there are many variants of sand-casting, some specialized to give greater precision, others modified to allow exceptional size, still others adapted to deal with specific materials.

Each member is characterized by a set of *attributes*. It has *material attributes*: the particular subset of materials to which it can be applied. It has *shape-creating attributes*: the classes of shapes

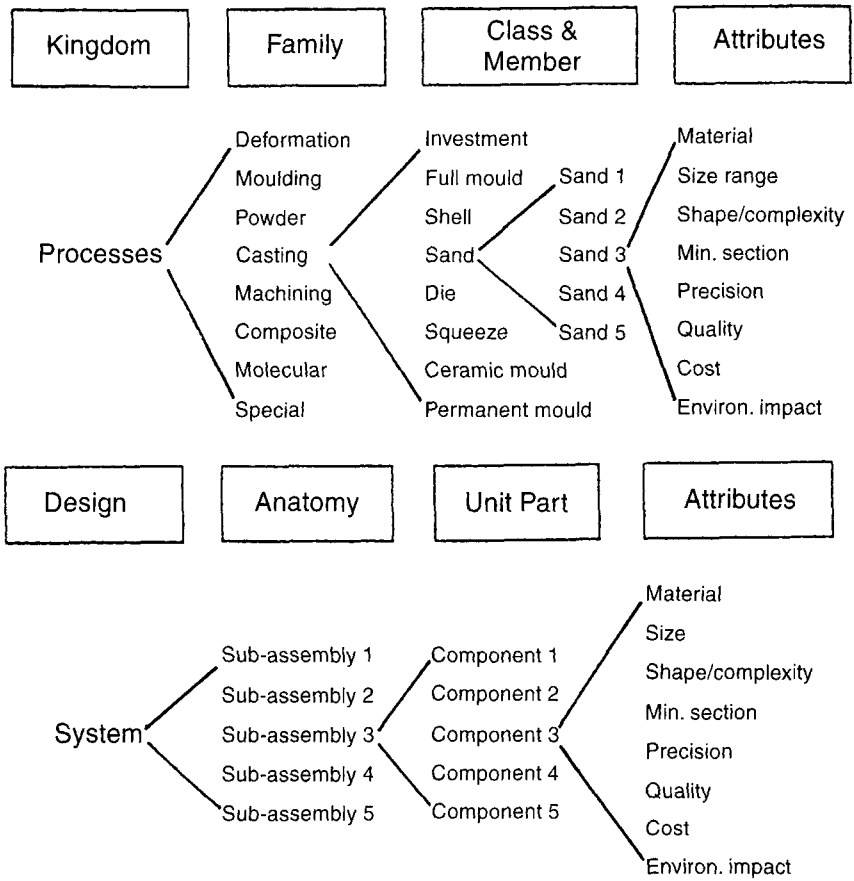

Fig. 11.26 Top: the taxonomy of the kingdom of process, and their attributes; bottom: the design of a component defines a required attribute profile. Process selection involves matching the two.

it can make. It has *physical attributes* which relate to the size, precision, finish and quality of its product. It has attributes which relate to the *economics* of its use: its capital cost and running cost, the speed with which it can be set up and operated, the efficiency of material usage. And it has attributes which relate to its impact on the environment: its eco-cost, so to speak.

Process selection is the action of matching process attributes to the attributes required by the design (Figure 11.26, bottom half). The anatomy of a design can be decomposed into *sub-assemblies*; these can be subdivided into *components*; and components have *attributes*, specified by the designer, some relating to material, some to shape, some physical, some economic. The problem, then, is that of matching the attribute-profiles of available processes to that specified by the design.

11.4 Systematic process selection

You need a process to shape a given material to a specified shape and size, and with a given precision. How, from among the huge number of possible processes, are you to choose it? Here is the strategy. The steps parallel those for selecting a material. In four lines:

(a) consider *all processes to be candidates* until shown to be otherwise;
(b) *screen* them, eliminating those which lack the attributes demanded by the design;
(c) *rank* those which remain, using relative cost as the criterion;
(d) *seek supporting information* for the top candidates in the list.

Figure 11.27 says the important things. Start with an open mind: initially, *all* processes are options. The design specifies a material a shape, a precision, a batch size, and perhaps more. The first step — that of *screening* — eliminates the process which cannot meet these requirements. It is done by comparing the attributes specified by the design (material, for instance, or shape or precision) with the attributes of processes, using hard copy or computer-generated Process Selection Charts described in a moment. Here, as always, decisions must be moderated by common sense: some design requirements are absolute, resulting in rejection, others can be achieved by constructing process-chains. As an example, if a process cannot cope with a *material* it must be rejected, but if its *precision* is inadequate, this can be overcome by calling on a secondary process such as machining.

Screening gives the processes which could meet the design requirements. The next step is to *rank* them using economic criteria. There are two ways of doing it. Each process is associated with an 'economic batch size-range' or EBS: it is the range over which that process is found to be cheaper than competing processes. The design specifies a batch size. Processes with an EBS which corresponds to the desired batch size are put at the top of the list. It is not the best way of ranking, but it is quick and simple.

Better is to rank by *relative cost*. Cost, early in the design, can only be estimated in an approximate way, but the cost differences between alternative process routes are often so large that the estimate allows meaningful ranking. The cost of making a component is the sum of the costs of the resources consumed in its production. These resources include materials, capital, time, energy, space and information. It is feasible to associate approximate values of these with a given process, allowing the relative cost of competing processes to be estimated.

Screening and ranking reduce the kingdom of processes to a small subset of potential candidates. We now need *supporting information*. What is known about each candidate? Has it been used before to make components like the one you want? What is its family history? Has it got hidden faults, character defects, so to speak? Such information is found in handbooks, in the data sheets

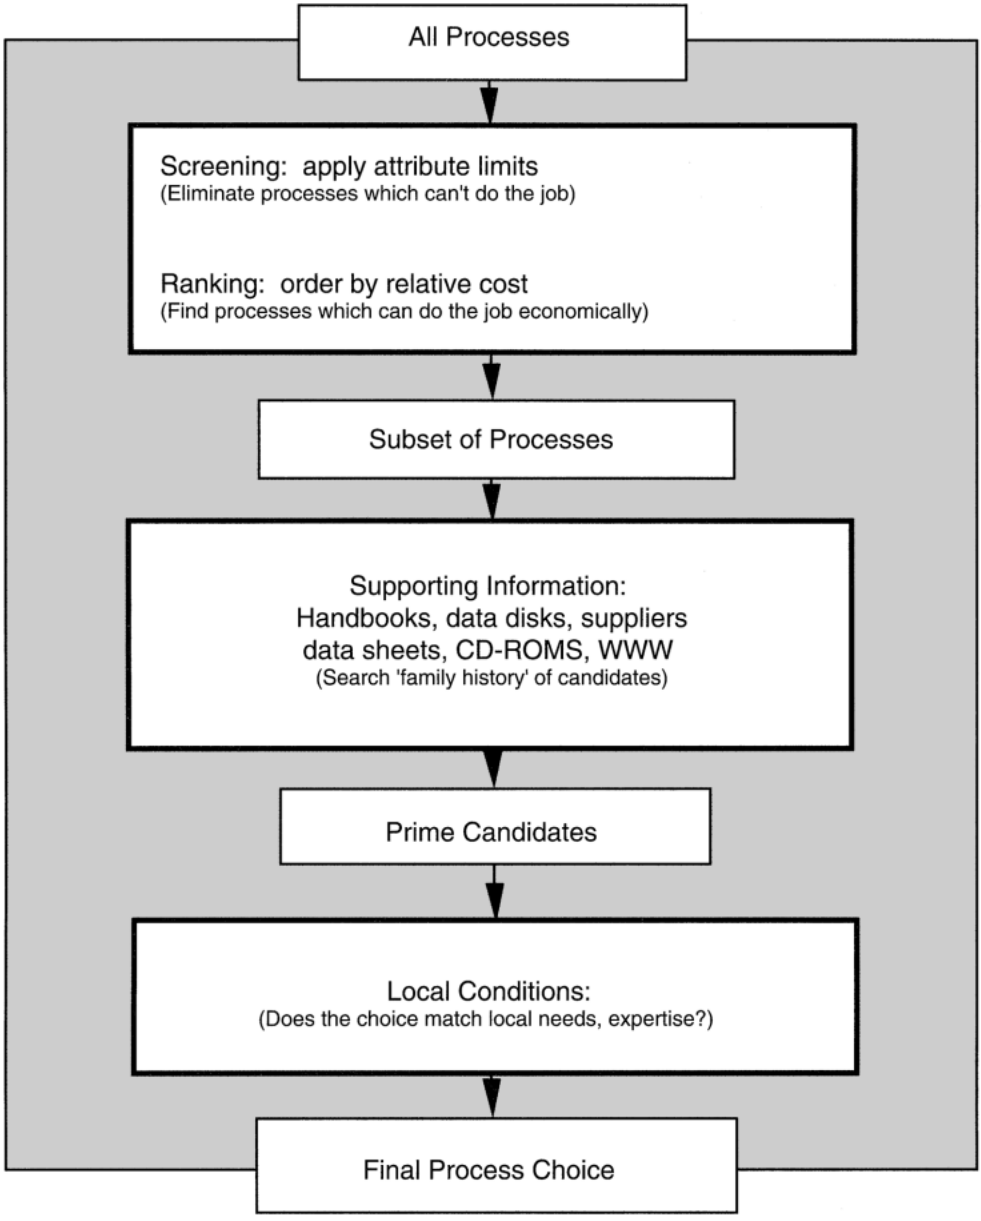
All Processes
Screening: apply attribute limits
(Eliminate processes which can't do the job)
Ranking: order by relative cost
(Find processes which can do the job economically)
Subset of Processes
Supporting Information:
Handbooks, data disks, suppliers
data sheets, CD-ROMS, WWW
(Search 'family history' of candidates)
Prime Candidates
Local Conditions:
(Does the choice match local needs, expertise?)
Final Process Choice

Fig. 11.27 A flow chart of the procedure for process selection. It parallels that for material selection.

of suppliers of process equipment and in documented case studies which, increasingly, appear in electronic format on CD or the World Wide Web.

This is as far as a general strategy can go. In reality there is one more step: it is to examine whether local conditions modify the choice. Available equipment and expertise for one class of process and lack of them for another can, for obvious reasons, bias the selection. But one should be aware that the unbiased choice might, in the long run, be better. That is the value of a systematic

strategy such as this one: it reveals the options and their relative merit. The final choice is up to the user.

11.5 Screening: process selection diagrams

Screening using hard copy diagrams

How do you find the processes which can form a given material to a given size, shape, and precision? First eliminate all processes which cannot handle the material; then seek the subset of these which can handle the size, create the shape and achieve the precision you want. Progress can be made by using the hard copy charts shown first in this section. Greater resolution is possible with computer-aided process selection software, described after that.

The axes of a process selection chart are measures of two of the attributes — precision and surface finish, for example. Figure 11.28 is a schematic of such a chart. The horizontal axis is the RMS surface roughness, plotted on a logarithmic scale, running from $10^{-3}\,\mu m$ to $100\,\mu m$. The vertical axis is the tolerance ranging from $\pm 10^{-4}$ mm to ± 10 mm. Each process occupies a particular area of the chart: it is capable of making components in a given range of tolerance and of roughness. Conventional casting processes, for instance, can make components with a tolerance in the range ± 0.1 to ± 10 mm (depending on process and size) with a roughness ranging from 5 to $100\,\mu m$; precision casting can improve both by a factor of 10. Machining adds precision: it extends the range down to $T = \pm 10^{-3}$ mm and $R = 0.01\,\mu m$. Polymer forming processes give high surface finish but limited tolerance. Lapping and polishing allow the highest precision and finish of all.

Selection is achieved by superimposing on the chart the envelope of attributes specified by the design, as shown in the figure. Sometimes the design sets upper and lower limits on process attributes (here: T and R), defining a closed box like that of Search area 1 of the figure. Sometimes, instead, it prescribes upper limits only, as in Search area 2. The processes which lie within or are bounded by the shaded search envelope are candidates; they are the initial shortlist. The procedure is repeated using similar charts displaying other attributes, narrowing the shortlist to a final small subset of processes capable of achieving the design goal.

There are some obvious difficulties. Process attributes can be hard to quantify: 'shape', for example, is not easy to define and measure. Certain processes have evolved to deal with special needs and do not naturally appear on any of the charts. Despite this, the procedure has the merits that it introduces a systematic element into process selection, and it forms the basis of a more sophisticated computer-based approach, described in a moment.

Material compatibility. The match between process and material is established by the link to material class and by the use of the material compatibility chart of Figure 11.29. Its axes are melting point and hardness. The melting point imposes limits on the processing of materials by conventional casting methods. Low melting point metals can be cast by any one of many techniques. For those which melt above 2000 K, conventional casting methods are no longer viable, and special techniques such as electron-beam melting must be used. Similarly, the yield strength or hardness of a material imposes limitations on the choice of deformation and machining processes. Forging and rolling pressures are proportional to the flow strength, and the heat generated during machining, which limits tool life, also scales with the ultimate strength or hardness. Generally speaking, deformation processing is limited to materials with hardness values below 3 GPa. Other manufacturing methods exist which are not limited either by melting point or by hardness. Examples are: powder methods, CVD and evaporation techniques, and electro-forming.

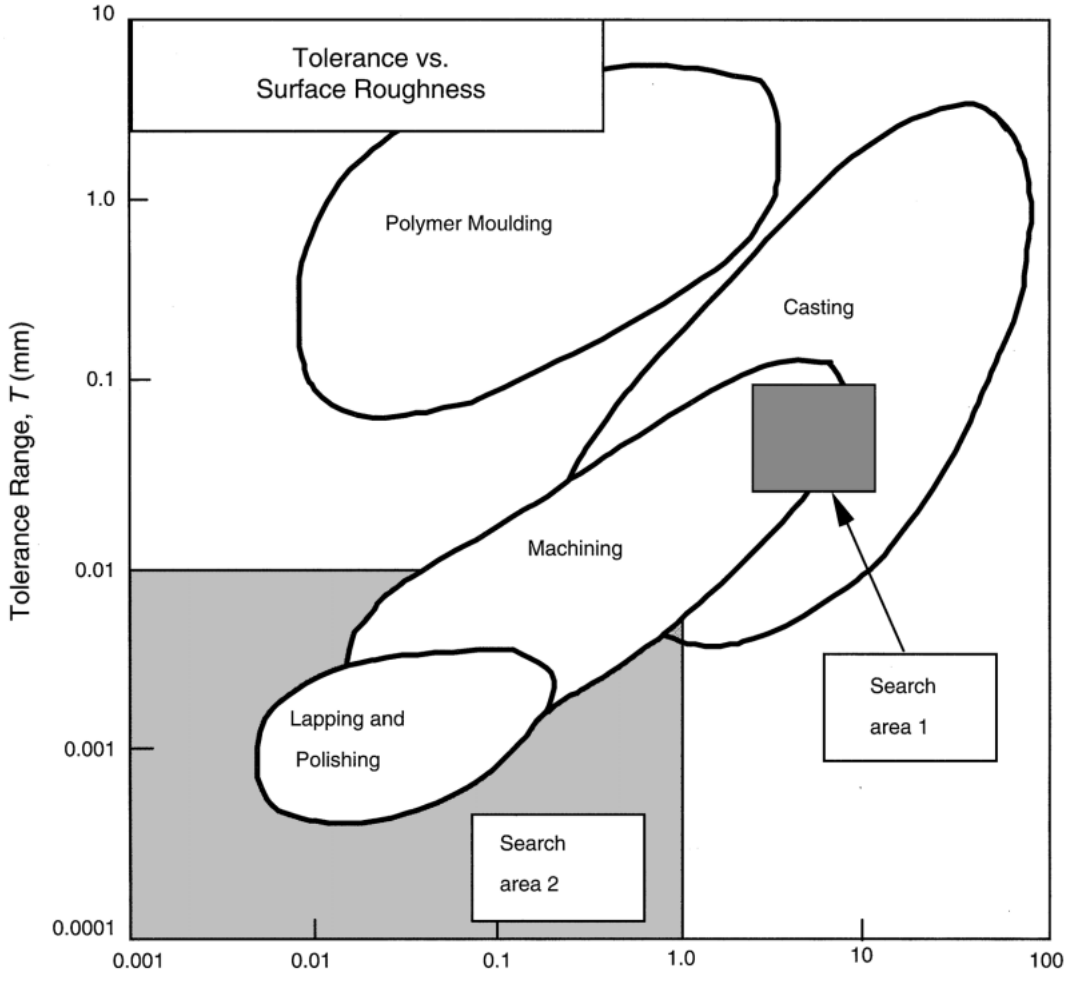

Fig. 11.28 A schematic illustrating the idea of a process-selection chart. The charts have process attributes as axes; a given process occupies a characteristic field. A design demands a certain set of processes attributes, isolating a box ('Search Area 1') or a sub-field ('Search Area 2') of the chart. Processes which overlap the search areas become candidates for selection.

Figure 11.29 presents this information in graphical form. In reality, only part of the space covered by the axes is accessible: it is the region between the two heavy lines. The hardness and melting point of materials are not independent properties: low melting point materials tend to be soft (polymers and lead, for instance); high melting point materials are hard (diamond is the extreme example). This information is captured by the equation

$$0.03 < \frac{H\Omega}{kT_m} < 20 \qquad (11.\ldots)$$

where Ω is the atomic or molecular volume and k is Boltzmann's Constant (1.38×10^{-26} J/K). It is this equation which defines the two lines.

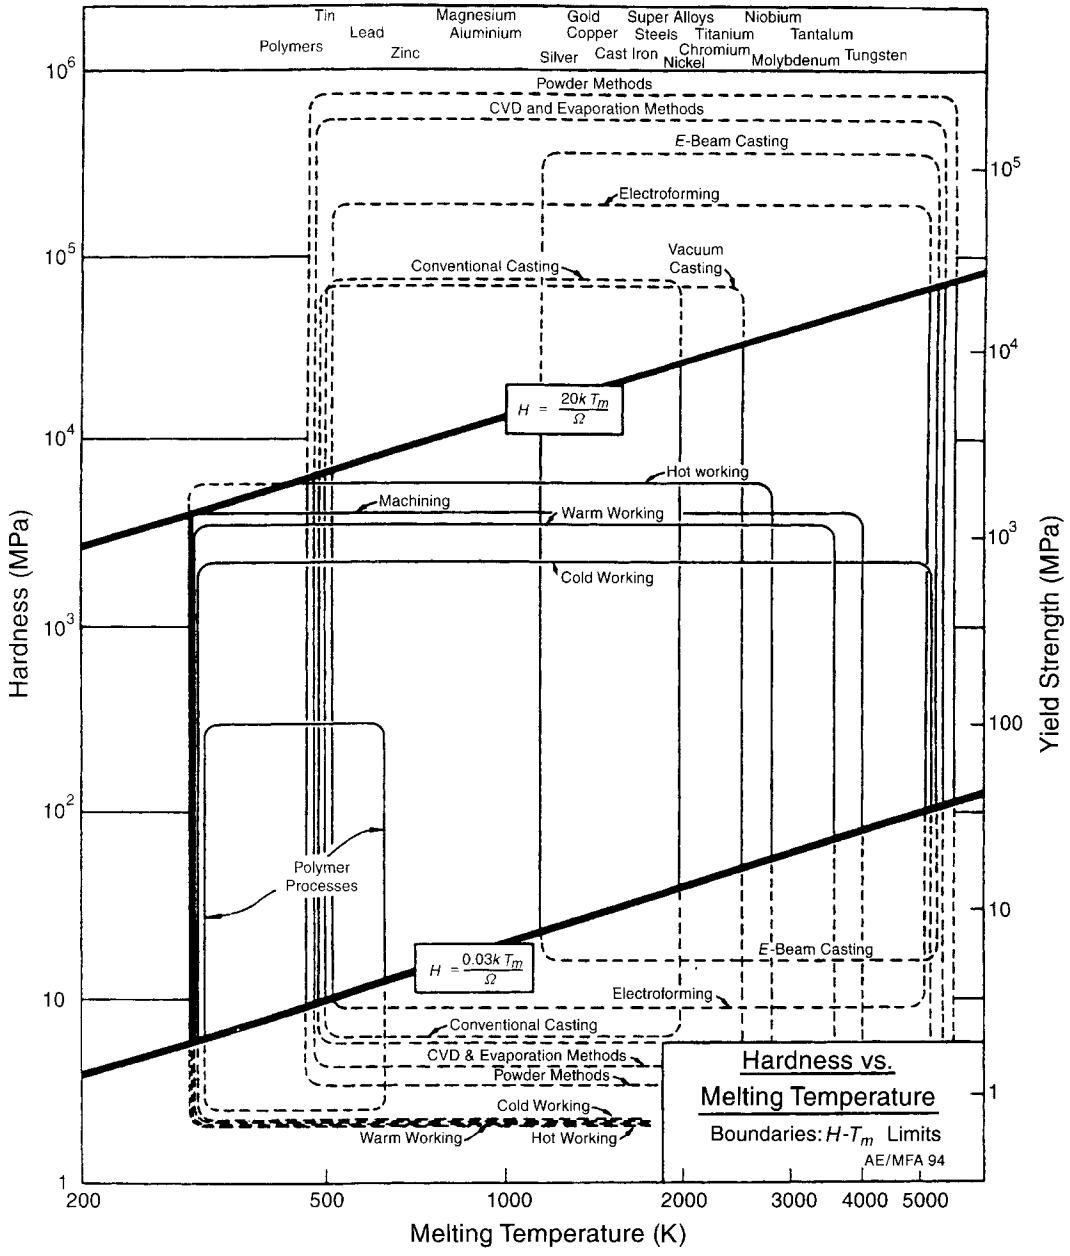

Fig. 11.29 The hardness–melting point chart.

Complexity, and the size–shape chart. Shape and complexity are the most difficult attributes to quantify. Pause for a moment to consider one way of quantifying complexity because it illustrates the nature of the difficulty. It is the idea of characterizing shape and complexity by *information content*. It has two aspects. The obvious one is the number n of independent dimensions which must be specified to describe the shape: for a sphere, it is 1 (the radius), for a cylinder, 2; for

a tube, 3. A complex casting might have 100 or more specified dimensions. Second there is the precision with which these dimensions are specified. A sphere of radius $r = 10\,\text{m} \pm 0.01\,\text{mm}$ is more 'complicated', in a manufacturing sense, than one with a radius $r = 10\,\text{mm} \pm 1\,\text{mm}$ because it is harder to make. Both aspects of complexity are captured by the information content

$$C = n \log_2 \left(\frac{\bar{\ell}}{\overline{\Delta\ell}} \right) \tag{11.2}$$

Here $\bar{\ell}$ is the average dimension and $\overline{\Delta\ell}$ is the mean tolerance (see Suh, 1990 for an extensive discussion). It looks as if it makes sense. The information content increases linearly with the number of dimensions, n, and logarithmically with the average relative precision $\bar{\ell}/\overline{\Delta\ell}$. The dimensions cease to have meaning if $\overline{\Delta\ell}$ equals $\bar{\ell}$ because the information content goes to zero.

So far, so good. But now compare a sphere (only one dimension) with a cylinder (with two). Spheres are hard to make, cylinders are not, even though they require twice as much information. Hollow spheres (two dimensions) are harder still, hollow tubes are easy. Information content does not relate directly to the way in which manufacturing processes actually work. Lathes are good at creating axisymmetric shapes (cylinders, tubes); rolling, drawing and extrusion are good at making prismatic ones (sheet, box-sections and the like). Add a single transverse feature and the processing, suddenly, becomes much more difficult. A measure of shape, if it is to be useful here, must recognize the capabilities and limitations of processes.

This directs our thinking towards axial symmetry, translational symmetry, uniformity of section and such like. As mentioned already, turning creates *axisymmetric* shapes; extrusion, drawing and rolling make *prismatic* shapes. Indexing gives shapes with *translational* or *rotational* symmetries, like a gear wheel. Sheet-forming processes make *flat* shapes (stamping) or *dished* shapes (drawing). Certain processes can make three-dimensional shapes, and among these, some can make hollow shapes, whereas others cannot. Figure 11.30 illustrates this classification scheme, building on those of Kusy (1976), Schey (1977) and Dargie *et al.* (1982). The shapes are arranged in the figure in such a way that complexity, defined here as the difficulty of making a shape, increases downwards and to the right.

Shape can be characterized in other ways. One, useful in process selection, is the aspect ratio, or what we call 'slenderness' S. Manufacturing processes vary widely in their capacity to make thin, slender sections. For our purposes, slenderness, S, is measured by the ratio t/ℓ where t is the minimum section and ℓ is the large dimension of the shape: for flat shapes, ℓ is about equal to $\sqrt{A}$ where A is the projected area normal to t. Thus

$$S = \frac{t}{\sqrt{A}} \tag{11.3}$$

Size is defined by the minimum and maximum volumes of which the process is capable. The volume, V, for uniform sections is, within a factor of 2, given by

$$V = At \tag{11.4}$$

Volume can be converted approximately to weight by using an 'average' material density of $5000\,\text{kg/m}^3$; most engineering materials have densities within a factor of 2 of this value. Polymers are the exception: their densities are all around $1000\,\text{kg/m}^3$.

The size–slenderness chart is shown in Figure 11.31. The horizontal axis is the slenderness, S; the vertical axis is the volume, V. Contours of A and t are shown as families of diagonal lines. Casting processes occupy a characteristic field of this space. Surface tension and heat-flow limit

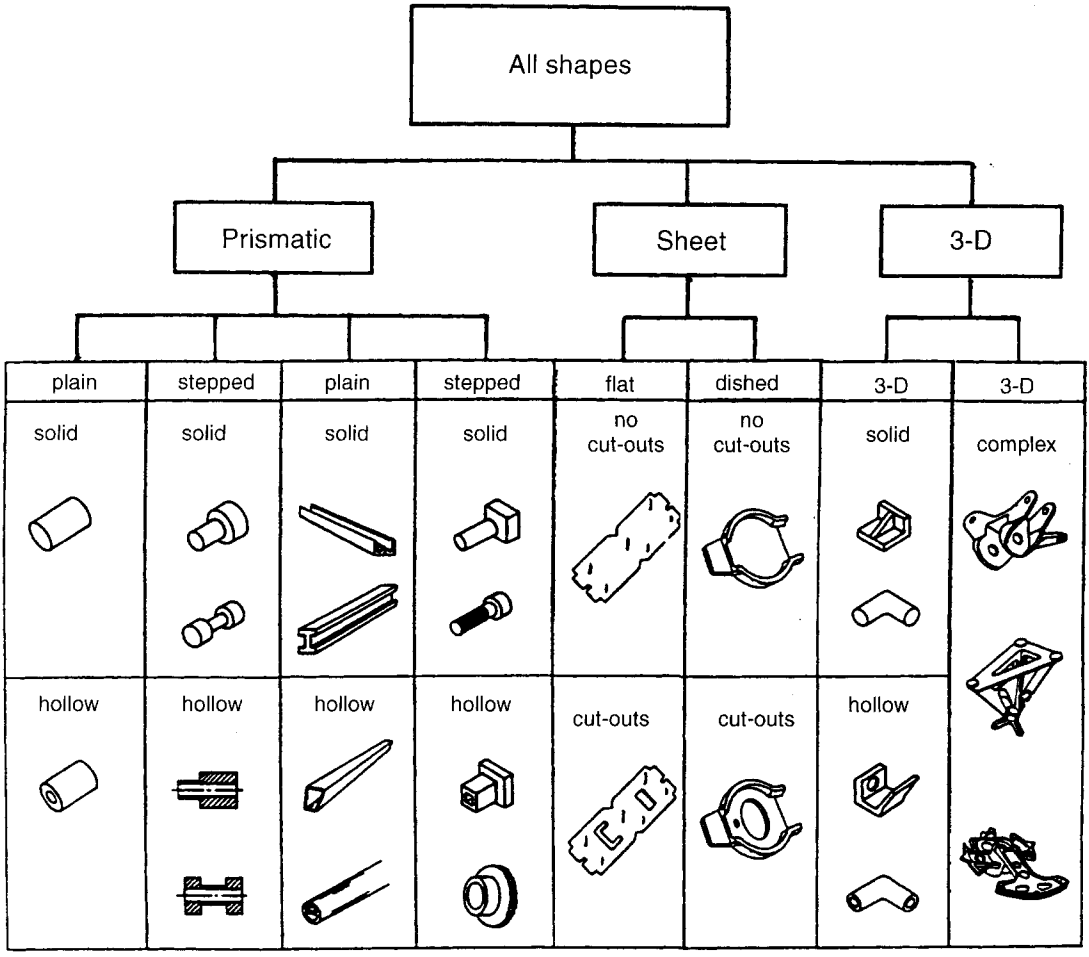

Fig. 11.30 A classification of shape that correlates with the capabilities of process classes.

the minimum section and the slenderness of gravity cast shapes. The range can be extended by applying a pressure, as in centrifugal casting and pressure die casting, or by preheating the mould. But there remain definite upper and lower limits to the size and shape achievable by casting. Deformation processes — cold, warm and hot — cover a wider range. Limits on forging-pressures set a lower limit on thickness and slenderness, but it is not nearly as severely as in casting. Sheet, wire and rod can be made in very great lengths — then the surface area becomes enormous. Machining creates slender shapes by removing unwanted material. Powder-forming methods occupy a smaller field, one already covered by casting and deformation shaping methods, but they can be used for ceramics and very hard metals which cannot be shaped in other ways. Polymer-forming methods — injection moulding, pressing, blow-moulding, etc. — share this regime. Special techniques, which include electro-forming, plasma-spraying, and various vapour-deposition methods, allow very slender shapes. Micro-fabrication technology, in the extreme lower part of the chart, refers to the newest techniques for sub-micron deposition and chemical or electron-beam milling. Joining extends the range further: fabrication allows almost unlimited size and complexity.

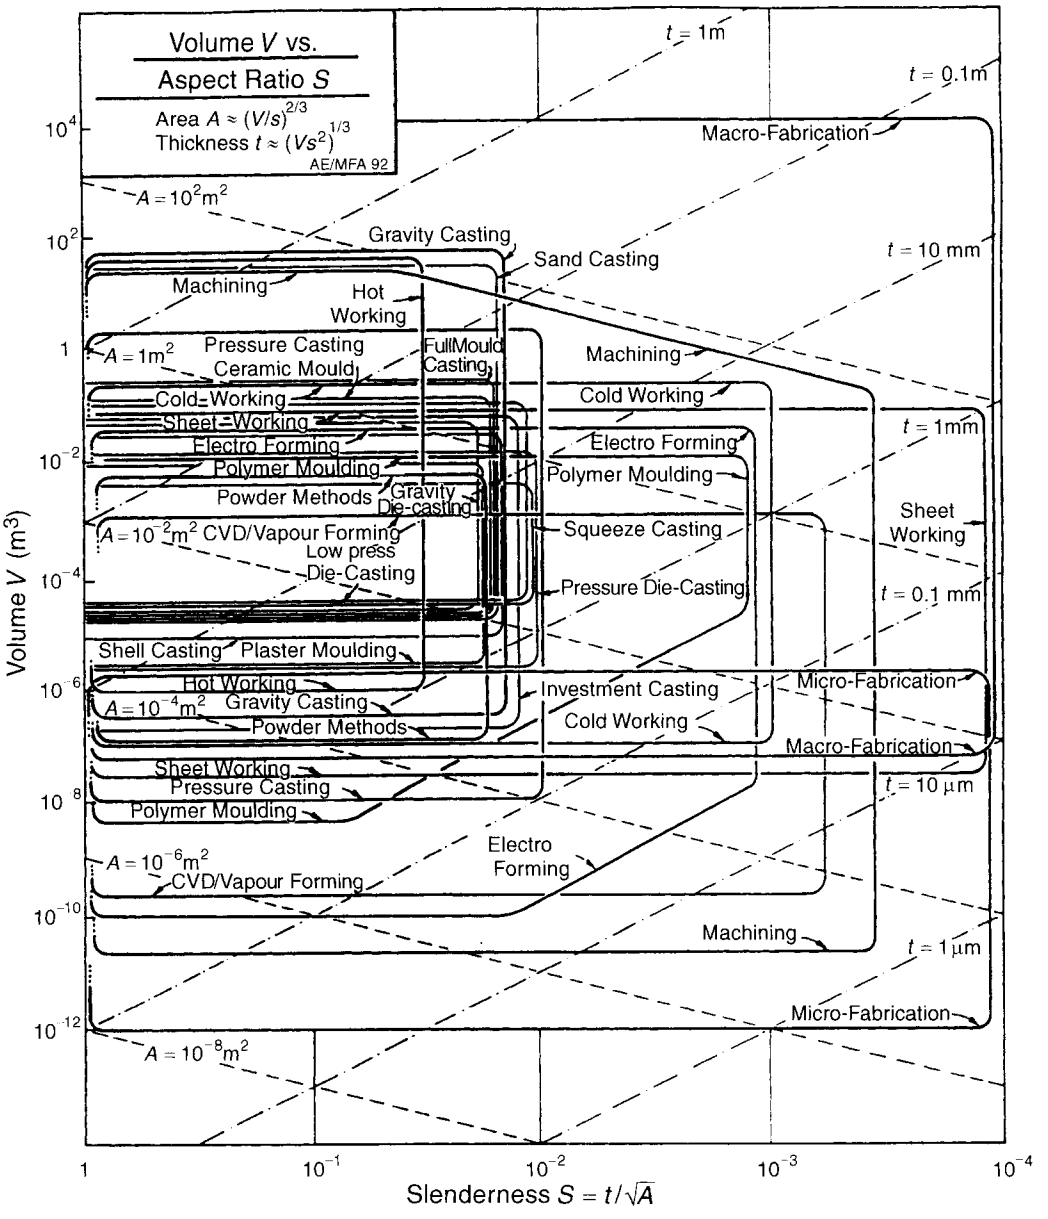

Fig. 11.31 The size–slenderness chart. Diagonal contours give approximate measures of area A and thickness t.

A real design demands certain specific values of S and V, or A and t. Given this information, a subset of possible processes can be read off. Examples are given in the next chapter.

The complexity level vs. *size chart.* Complexity is defined as *the presence of features such as holes, threads, undercuts, bosses, re-entrant shapes, etc., which cause manufacturing difficulty or require additional operations.* So if Figure 11.31 describes the basic shape, complexity describes the

additional extra features which are required to produce the final shape. For purposes of comparison, a scale of 1 to 5 is used with 1 indicating the simplest shapes and 5 the most complicated. Each process is given a rating for the maximum complexity of which it is capable corresponding to its proximity to the top left or bottom right shapes in Figure 11.30.

This information is plotted on the complexity level–size chart shown in Figure 11.32. Generally, deformation processes give shapes of limited complexity. Powder routes and composite forming methods are also limited compared with other methods. Polymer moulding does better. Casting processes offer the greatest complexity of all: a cast automobile cylinder block, for instance, is an extremely complicated object. Machining processes increase complexity by adding new features to a component. Fabrication extends the range of complexity to the highest level.

The tolerance–surface roughness chart. No process can shape a part *exactly* to a specified dimension. Some deviation Δx from a desired dimension x is permitted; it is referred to as the *tolerance*,

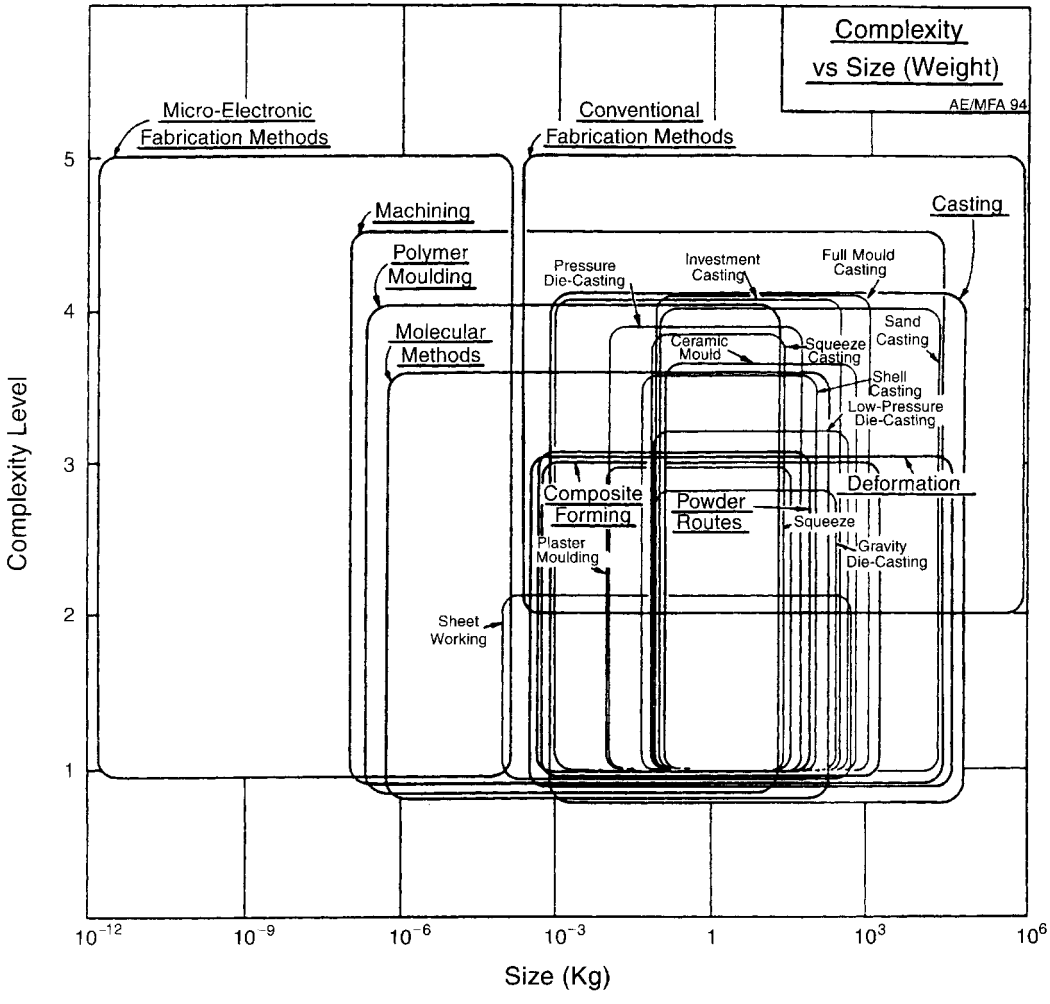

Fig. 11.32 The complexity–size chart.

Table 11.1 Levels of finish

Finish, μm	*Process*	*Typical application*
$R = 0.01$	Lapping	Mirrors
$R = 0.1$	Precision grind or lap	High quality bearings
$R = 0.2–0.5$	Precision grinding	Cylinders, pistons, cams, bearings
$R = 0.5–2$	Precision machining	Gears, ordinary machine parts
$R = 2–10$	Machining	Light-loaded bearings, Non-critical components
$R = 3–100$	Unfinished castings	Non-bearing surfaces

T, and is specified as $x = 100 \pm 0.1$ mm, or as $x = 50^{+0.01}_{-0.001}$ mm. Closely related to this is the *surface roughness* R, measured by the root-mean-square amplitude of the irregularities on the surface. It is specified as $R < 100\,\mu$m (the rough surface of a sand casting) or $R < 0.01\,\mu$m (a lapped surface; Table 11.1).

Manufacturing processes vary in the levels of tolerance and roughness they can achieve economically. Achievable tolerances and roughnesses are shown in Figure 11.33. The tolerance is obviously greater than $2R$ (shaded band): indeed, since R is the root-mean-square roughness, the peak roughness is more like $5R$. Real processes give tolerances which range from about $10R$ to $1000R$. Sand casting gives rough surfaces; casting into metal dies gives a better finish. Moulded polymers inherit the finish of the moulds and thus can be very smooth, but tolerances better than ± 0.2 mm are seldom possible because of internal stresses left by moulding and because polymers creep in service. Machining, capable of high dimensional accuracy and smooth surface finish, is commonly used after casting or deformation processing to bring the tolerance or finish to the desired level. Metals and ceramics can be surface-ground and lapped to a high tolerance and smoothness: a large telescope reflector has a tolerance approaching 5 μm over a dimension of a metre or more, and a roughness of about 1/100 of this. But such precision and finish are expensive: processing costs increase almost exponentially as the requirements for tolerance and surface finish are made more severe. The chart shows contours of relative cost: an increase in precision corresponding to the separation of two neighbouring contours gives an increase in cost, for a given process, of a factor of two. It is an expensive mistake to overspecify precision.

Achievable tolerances depend, of course, on dimensions (those given here apply to a 25 mm dimension) and on material. However, for our purposes, typical ranges of tolerance and surface finish are sufficient and discriminate clearly between various processes.

Use of hard copy process selection charts. The charts presented here provide an overview: an initial at-a-glance graphical comparison of the capabilities of various process classes. In a given selection exercise they are not all equally useful: sometimes one is discriminating, another not — it depends on the design requirements. They should not be used blindly, but used to give guidance in selection and engender a feel for the capabilities and limitations of various process types, remembering that some attributes (precision, for instance) can be added later by using secondary processes. That is as far as one can go with hard-copy charts. The number of processes which can be presented on them is obviously limited and the resolution is poor because so many of them overlap. But the procedure lends itself well to computer implementation, overcoming these deficiencies.

Computer-aided screening

If process attributes are stored in a database with an appropriate user-interface, selection charts can be created and selection boxes manipulated with much greater freedom. The Cambridge Process

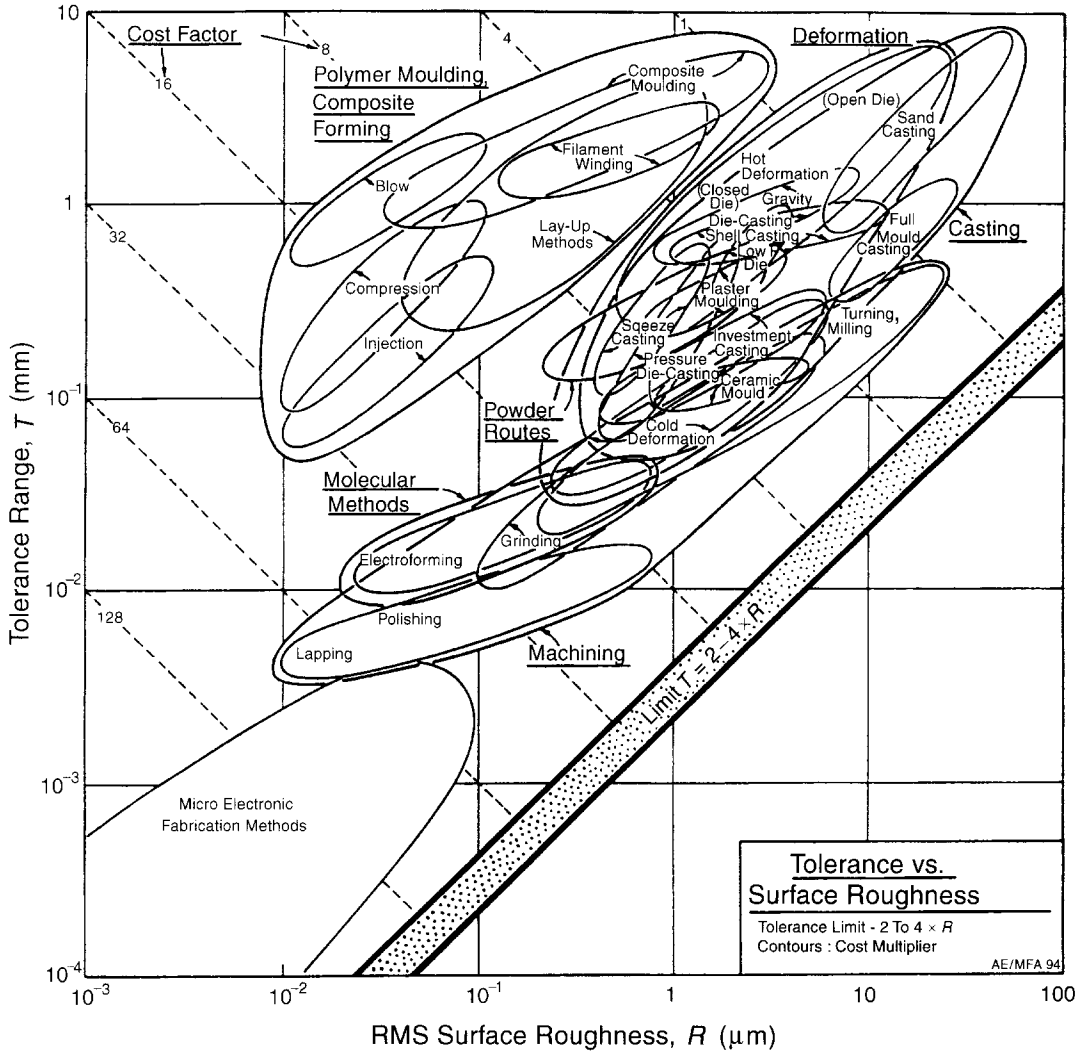

Fig. 11.33 The tolerance–roughness chart.

Selector *(CPS)* is an example of such a system. The way it works is described here; examples of its use are given in Chapter 12. The database contains records, each describing the attributes of a single process. Figure 11.34 shows a typical record: it is that for a particular member of the sand casting class. A schematic indicates how the process works; it is supported by a short description. This is followed by a listing of attributes: the material capability, the attributes relating to shape and physical characteristics, and those which describe economic parameters; the record also contains a brief description of typical uses, references and notes. All the numeric attributes are stored as ranges, indicating the range of capability of the process. The record concludes with a set of references from which the data were drawn and which provide intelligence information, essential in reaching a final selection.

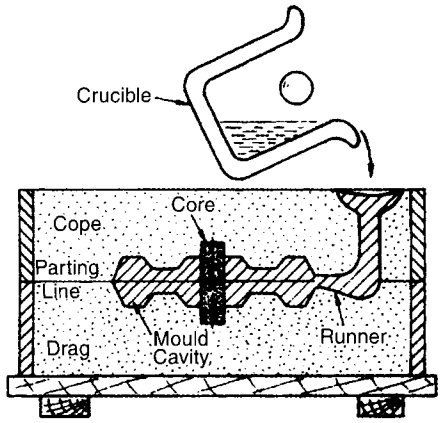

Fig. 11.34 A typical record from a computer-based process selector. It is that for a member of the casting family: CO_2/Silicate sand casting.

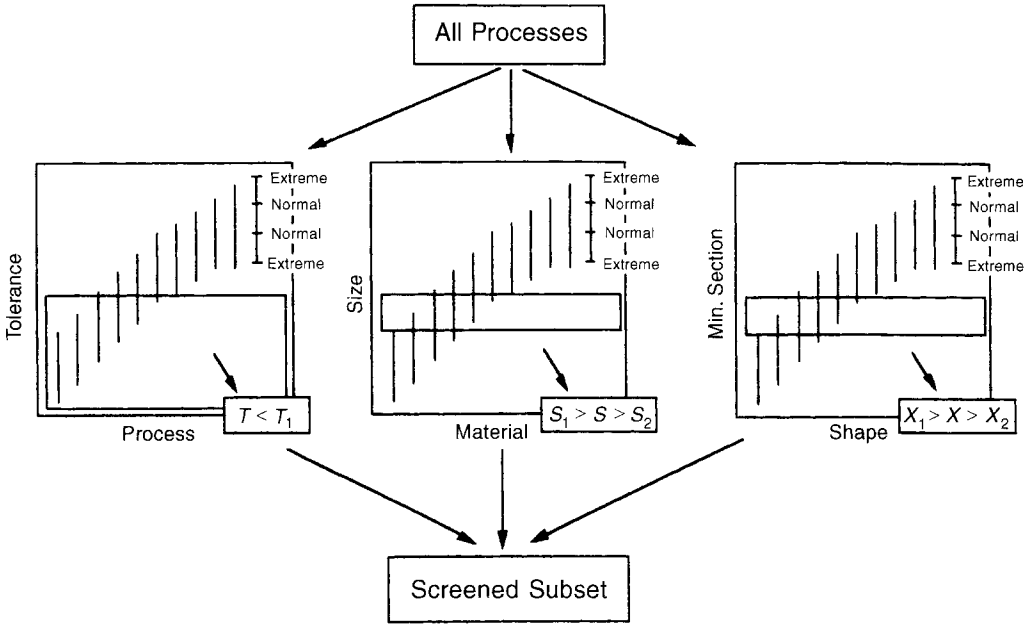

Fig. 11.35 Computer-based screening. The attributes of processes are plotted as bar-charts, isolating processes which can handle a given material class, create a given class of shapes, and meets the design requirements on size, tolerance and minimum section.

The starting point, as in Figure 11.27, is the idea that all processes are potential candidates until shown otherwise. A shortlist of candidates is extracted in two steps: screening to eliminate processes which cannot meet the design specification, and ranking to order the survivors by economic criteria.

A typical three stage screening takes the form shown in Figure 11.35. It shows three bar-charts, on each of which a numeric property is plotted for a selected class property (process class, material

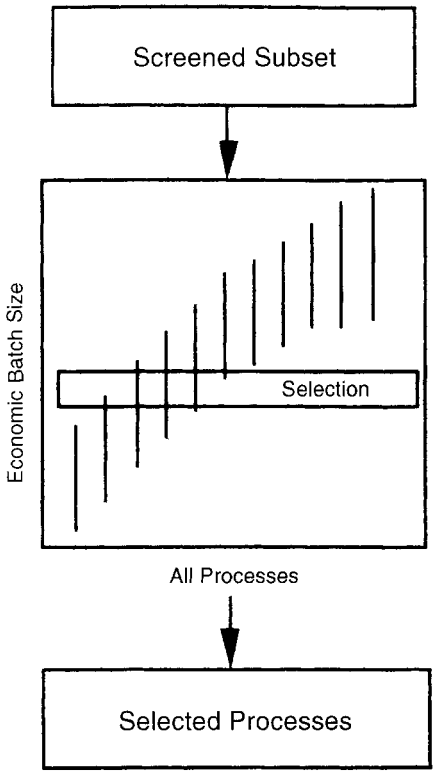

Fig. 11.36 Ranking by economic criteria, here, the economic batch size.

class and shape class). All processes with the selected class attributes appear on the charts. The processes are sorted in order of ascending value of the numeric property which is plotted as a bar to show its range. The left-hand chart selects a process of a given type ('primary', for example) which offers a tolerance better than $\pm T_1$ mm. The second specializes this to those which can shape a chosen class of material ('thermoplastic polymers', for instance) with a size between S_1 and S_2 kg. The third isolates the subset of these which are able to create a given shape (such as '3-D solid, parallel features') with a minimum section thickness as small as X_1 mm. Further stages can be added. The selection is made by placing a selection box onto each chart, identifying the range of tolerance, size, minimum section and so forth specified by the design. The effect is to eliminate the processes which cannot meet the specifications.

The next step is to rank the survivors by economic criteria (Figure 11.36). To do this we need to examine process cost.

11.6 Ranking: process cost

Part of the cost of a component is that of the material of which it is made. The rest is the cost of manufacture, that is, of forming it to a shape, and then of joining it to the other components to give the finished product. Before turning to details, there are three common-sense rules for minimizing cost which the designer should bear in mind.

Keep things standard. If someone already makes the part you want, it will almost certainly be cheaper to buy it than to make it. If nobody does, then it is cheaper to design it to be made from standard stock (sheet, rod, tube) than from non-standard shapes, or special castings or forgings. Try to use standard materials, and as few of them as possible: it reduces inventory costs and the range of tooling the manufacturer needs.

Keep things simple. If a part has to be machined, it will have to be clamped; the cost increases with the number of times it will have to be re-jigged or re-oriented, specially if special tools are necessary. If a part is to be welded or brazed, the welder must be able to reach it with his torch and still see what he is doing. If it is to be cast or moulded or forged, it should be remembered that high (and expensive) pressures are required to make fluids flow into narrow channels, and that re-entrant shapes greatly complicate mould design. All this is pretty obvious, but easily overlooked. Think of making the part yourself: will it be awkward? Could slight re-design make it less awkward?

Do not specify more performance than is needed. Performance must be paid for. High-strength metals are more heavily alloyed with expensive additions; high-performance polymers are chemically more complex; high-performance ceramics require greater quality control in their manufacture. All of these increase material costs. In addition, high-strength materials are hard to fabricate. The forming pressures (whether for a metal or a polymer) are higher; tool wear is greater; ductility is usually less so that deformation processing can be difficult or impossible. This can mean that new processing routes must be used: investment casting or powder forming instead of conventional casting and mechanical working; more expensive moulding equipment operating at higher temperatures and pressures, and so on. The better performance of the high-strength material must be paid for, not only in greater material cost but also in the higher cost of processing. Finally, there are the questions of tolerance and roughness. The 'cost' contours of Figure 11.33 give warning: cost rises exponentially with precision and surface finish. It is an expensive mistake to specify tighter tolerance or smoother surfaces than are necessary. The message is clear. Performance costs money. Do not over specify it.

To make further progress, we must examine the contributions to process costs, and their origins.

Economic criteria for selection. If you have to sharpen a pencil, you can do it with a knife. If, instead, you had to sharpen a thousand pencils, it would pay to buy an electric 'mechanized' sharpener. And if you had to sharpen a million, you might wish to equip yourself with an automatic feeding, gripping and sharpening system. To cope with pencils of different length and diameter, you could go further and devise a microprocessor-controlled system with sensors to measure pencil dimensions, sharpening pressure and so on; that is, you create a system with 'intelligence' which can recognize and adapt to pencil size. The choice of process, then, depends on the number of pencils you wish to sharpen, that is, on the *batch size*. The best choice is that which costs least, per pencil sharpened.

Figure 11.37 is a schematic of how the cost of sharpening a pencil might vary with batch size. A knife does not cost much but it is slow, so the labour cost is high. The other processes involve progressively greater capital investment but do the job more quickly, reducing labour costs. The balance between capital cost and rate gives the shape of the curves. In this figure the best choice is the lowest curve — a knife for up to 100 pencils; mechanization for 10^2 to 10^4, an automatic system for 10^4 to 10^6, and so on.

Economic batch size. Modelling cost may sound easy but it is not. Process cost depends on a large number of independent variables, not all within the control of the modeller. Cost modelling is described in the next section, but — given the disheartening implications of the last sentence — it is

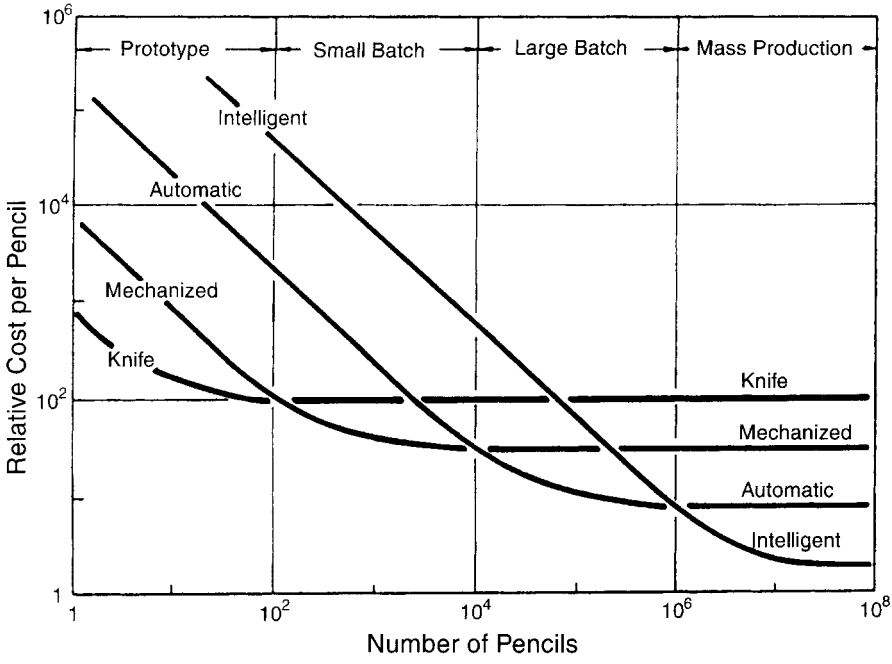

Fig. 11.37 The cost of sharpening a pencil, plotted against batch size, for four processes. The curves all have the form of equation (11.7).

comforting to have an easy, if approximate, way out. The influence of many of the inputs to the cost of a process are captured by a single attribute: the *economic batch size*. A process with an economic batch size with the range $B_1 - B_2$ is one which is found by experience to be competitive in cost when the output lies in that range. The economic batch size is commonly cited for processes. The easy way to introduce economy into the selection is to rank candidate processes by economic batch size and retain those which are economic in the range you want, as illustrated by Figure 11.36. But do not harbour false illusions: many variables cannot be rolled into one without loss of discrimination. It is better to develop a cost model.

Cost modelling. The manufacture of a component consumes resources (Table 11.2). The process cost is the sum of the costs of the resources it consumes. This *resource-based* approach to cost analysis is particularly helpful at the broad level with which we are concerned here since all processes, no matter how diverse, consume the resources listed in the table. Thus the cost of a component of mass m entails the cost C_m (\$/kg) of the *material* of which it is made, and it involves the cost of *dedicated tooling*, C_t, which must be amortized by the *batch size*, n. In addition, it requires *time*, chargeable at an *overhead rate* $\dot{C}_L$ (thus with units of \$/h or equivalent), *power* $\dot{P}$ (kW) at an *energy cost* C_e (\$/kWh), and it requires *space* of area A, incurring a *rental cost* of $\dot{C}_s$ (\$/m^{2}h). The cost equation takes the form

$$\begin{matrix} & \text{Material} & \text{Tooling} & \text{Time} & \text{Energy} & \text{Space} \\ C = & [mC_m] & + \left[\frac{C_t}{n}\right] & + \left[\frac{\dot{C}_L}{\dot{n}}\right] & + \left[\frac{\dot{P}C_e}{\dot{n}}\right] & + \left[\frac{A\dot{C}_s}{\dot{n}}\right] \end{matrix} \qquad (11.5)$$

where $\dot{n}$, the *batch rate*, is the number of units produced per hour.

Table 11.2 The resources consumed in production

Resource		*Symbol*	*Unit*
Materials:	inc. consumables	C_m	\$/kg
Capital:	of equipment	C_c	\$
	cost of tooling	C_t	\$
Time:	overhead rate	$\dot{C}_L$	\$/hr
Energy:	power	$\dot{P}$	kW
	cost of energy	C_e	\$/kW h
Space:	area	A	m^2
	cost of space	$\dot{C}_s$	$\$/m^2h$
Information:	R & D	C_i	\$/yr
	royalty payments		

Where has the *capital cost* C_c of the equipment (as opposed to tooling) gone? A given piece of equipment — a press, for example — is commonly used to make more than one product. It is then usual to convert the capital cost of non-dedicated equipment, and the cost of borrowing the capital itself, into an overhead by dividing it by a capital write-off time, t_c (5 years, say) over which it is to be recovered. Thus the overhead rate becomes

$$\underset{\text{Basic OH rate}}{\frac{\dot{C}_L}{\dot{n}}} = \frac{1}{\dot{n}}\left\{\underset{\text{Basic OH rate}}{[\dot{C}_{Lo}]} + \underset{\text{Capital write-off}}{\left[\frac{C_c}{Lt_c}\right]}\right\} \tag{11.6}$$

where $\dot{C}_{Lo}$ is the *basic overhead rate* (labour, etc.) and L is the *load factor* (the fraction of time over which the equipment is productively used).

A detailed analysis breaks cost down further, detailing the contributions of scrap, administration, maintenance, the cost of capital (the interest that must be paid, or could have been earned, on the capital tied up in the equipment) and so on — real cost models can become very complex. Let us, instead, simplify. The terms can be assembled into three groups:

$$C = \underset{\text{Materials}}{[mC_m]} + \underset{\text{Tooling}}{\frac{1}{n}[C_t]} + \frac{1}{\dot{n}}\left[\underset{\text{Time}}{\dot{C}_{Lo}} + \underset{\text{Capital}}{\frac{C_c}{Lt_c}} + \underset{\text{Energy}}{\dot{P}C_e} + \underset{\text{Space}}{AC_s}\right]$$

We merge the terms in the final bracket into a single 'gross overhead', $\dot{C}_{L,\text{gross}}$, allowing the equation to be written

$$C = \underset{\text{Materials}}{[mC_m]} + \underset{\text{Dedicated cost/unit}}{\frac{1}{n}[C_t]} + \underset{\text{Gross overhead/unit}}{\frac{1}{\dot{n}}[\dot{C}_{L,\text{gross}}]} \tag{11.7}$$

The equation really says: cost has three types of contributions — one which is independent of batch size and rate, one which varies as the reciprocal of the batch size (n^{-1}), and one which varies as the reciprocal of the batch rate ($\dot{n}^{-1}$). The first — the 'material' costs — includes also material consumed in manufacture. The second — the dedicated capital investment — contains the cost of tooling, dies, jigs and moulds. The last term — the one dependent on time — includes the 'direct' cost of the machine operator plus the 'indirect' or 'overhead' cost associated with administration, maintenance, safety, and so forth. It is sometimes difficult to decide precisely how costs should be

assigned between these headings; different companies do it in different ways. But the general point is clear: *material* plus *dedicated capital costs* plus *gross overhead.*

The equation describes a set of curves, one for each process. Each has the shape of the pencil-sharpening curves of Figure 11.37. Figure 11.38 illustrates a second example: the manufacture of an aluminium con-rod by two alternative processes: sand casting and die casting. Sand casting equipment is cheap but the process is slow. Die casting equipment costs much more but it is also much faster. Data for the terms in equation (11.7), for these two processes, are listed in Table 11.3: they show that the capital cost assigned to the die-casting equipment is greater by some 76 times that for sand casting, but that the process is 40 times faster. The material cost (1 unit) and the labour cost per hour (20 units) are, of course, the same for both. Figure 11.38 is a plot of equation (11.7), evaluated with this data for the two processes. The curves intersect at a batch size of 4000. Sand casting is the most economical process for batches less than this; die casting for batches which are larger. Note that, for small batches, the component cost is dominated by that of the process — the material cost hardly matters. But as the batch size grows, the contribution of the second term in the cost equation diminishes; and if the process is fast, the cost falls until it is typically about twice that of the material of which the component is made.

Technical cost modelling. Equation (11.7) is the first step in modelling cost. Greater predictive power is possible with technical cost models which exploit understanding of the way in which the design, the process and cost interact. The capital cost of equipment depends on size and degree of automation. Tooling cost and production rate depend on complexity. These and many other

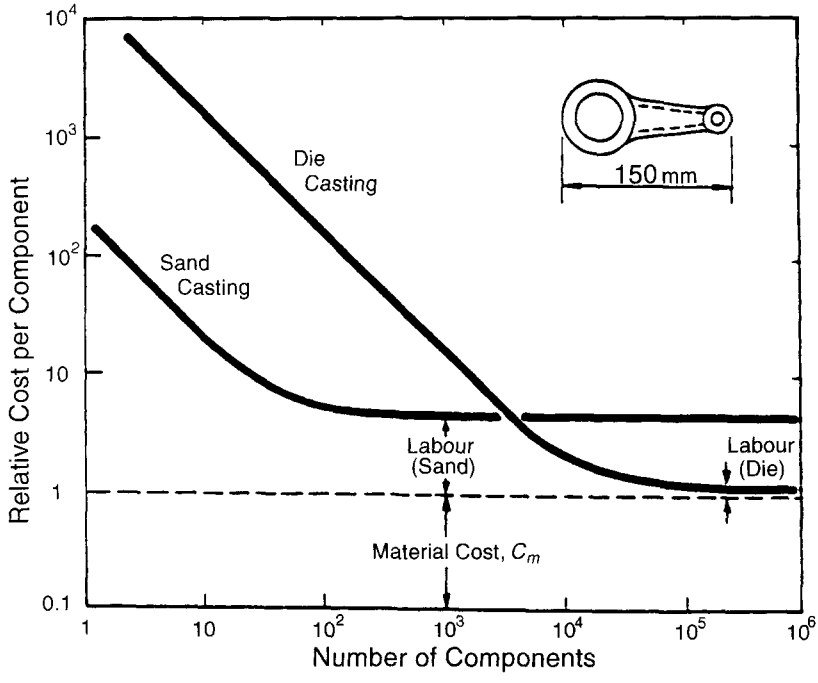

Fig. 11.38 The best choice of casting (or machining or forging) process depends on batch size. Sand casting requires cheap equipment but is labour intensive. Die casting requires more expensive equipment, but is faster. The data shown here are for an automobile connecting rod.

Table 11.3 Data for the cost equation

*Relative cost**	*Sand Casting*	*Die Casting*	*Comment*
Material, mC_m	1	1	Process independent
Basic overhead $C_{Lo}(h^{-1})$	20	20	
Capital write-off time t_c (yrs)	5	5	
Dedicated tool cost, C_t	210	16 000	Process dependent
Capital cost C_c	10 000	300 000	
Batch rate, $\dot{n}(h^{-1})$	5	200	

*All costs normalized to the material cost.

dependencies can be captured in theoretical or empirical formulae or look-up tables which can be built into the cost model, giving more resolution in ranking competing processes. For more advanced analyses the reader is referred to the literature listed in the Further reading section of this chapter.

11.7 Supporting information

Systematic screening and ranking based on attributes common to all processes gives a short list of candidates. We now need supporting information — details, case studies, experience, warnings, anything which helps form a final judgement. Where is it to be found?

Start with texts and handbooks — they don't help with systematic selection, but they *are* good on supporting information. They, and other sources are listed in Further reading. Bralla (1986) is particularly good and so is Schey (1977), although it is dated. Many texts and handbooks are specialized to a single class of process, giving more detail. Casting is an example: detailed intelligence is to be found in the ASM *Casing Design Handbook* (1962) and Clegg's (1991) *Precision Casting*.

Next look at the data sheets and design manuals available from the makers and suppliers of process equipment, and, often, from material suppliers. Leading suppliers exhibit at major conferences and exhibitions — these are a useful source of information for small and medium scale industries. Increasingly this sort of supplier-specific information is available on CD, allowing rapid access.

And then there is the World Wide Web. My old mother, were she still here, would have described it as 'a dog's dinner' meaning, I believe, that it contained everything from the best to the worst bits of today's menu. It is certainly mixed, but it *is* today's menu, and that has value. There are an increasing number of web sites which offer information on processes, the best of them very helpful. A selection is given in Chapter 13.

11.8 Summary and conclusions

A wide range of shaping and finishing processes is available to the design engineer. Each has certain characteristics, which, taken together, suit it to the forming of certain materials to certain shapes, but disqualify it for others. Faced with the choice, the designer has, in the past, relied on locally available expertise, or on common practice. Neither of them lead to innovation, nor are they well matched to current design methods. The structured, systematic approach of this chapter provides a way forward. It ensures that potentially interesting processes are not overlooked, and guides the user quickly to processes capable of making the desired shape.

The method parallels that for selection of material, using process selection charts to implement the procedure. The axes of the charts are process attributes: product size, shape, precision, and certain key material properties which influence shaping operations. A product design dictates a certain, known, combination of these attributes. The design requirements are plotted onto the charts, identifying a subset of possible processes.

There is, of course, much more to process selection than this. It is to be seen, rather, as a first systematic step, replacing a total reliance on local experience and past practice. The narrowing of choice helps considerably: it is now much easier to identify the right source for more expert knowledge and to ask of it the right questions. The final choice still depends on local economic and organizational factors which can only be decided on a case-by-case basis.

11.9 Further reading

Alexander, J.M., Brewer, R.C. and Rowe, G.W. (1987) *Manufacturing Technology, Vol. 2: Engineering Processes*, Ellis Horwood, Chichester.

ASM Casting Design Handbook (1962) American Society for Metals, Metals Park, Ohio USA.

Bolz, R.W. (1977) *Production Processes — the Productivity Handbook*, Conquest Publications, NC.

Bralla, J.G. (1986) *Handbook of Product Design for Manufacturing*, McGraw-Hill, New York.

Clegg, A.J. (1991) *Precision Casting Processes*, Pergamon Press, Oxford.

Crane, F.A.A. and Charles, J.A. (1984) *Selection and Use of Engineering Materials*, Butterworth, London, Chapter 13.

Dieter, G.E. (1983) *Engineering Design, A Materials and Processing Approach*, McGraw-Hill, New York, Chapter 7.

Edwards, L. and Endean, M., (eds) (1990) *Manufacturing with Materials*, Materials in Action Series, The Open University, Butterworths, London.

Esawi, A. (1994) 'Systematic Process Selection in Mechanical Design', PhD thesis, Cambridge University Engineering Department, Trumpington Street, Cambridge CB2 1PZ.

Esawi, A. and Ashby, M.F. (1998) 'Computer-based Selection of Manufacturing Processes, Part 1: methods and software; Part 2, case studies', Cambridge University Engineering Department Report TR 50, May 1997.

Esawi, A. and Ashby, M.F. (1998) 'Computer-based selection of manufacturing processes', *Journal of Engineering Manufacture*.

Farag, M.M. (1990) *Selection of Materials and Manufacturing Processes for Engineering Design*, Prentice-Hall, London.

Frost, H.J. and Ashby, M.F. (1982) *Deformation Mechanism Maps*, Pergamon Press, Oxford, Chapter 19, Section 111.6.

Kalpakjian, S. (1984) *Manufacturing Processes for Engineering Materials*, Addison Wesley, London.

Kusy, P.F. (1976) 'Plastic materials selection guide', SAE Technical Paper 760663, Sept. 1976.

Lascoe, O.D. (1988) *Handbook of Fabrication Processes*, ASM International, Metals Park, Columbus, Ohio.

Ludema, K.C., Caddell, R.M. and Atkins, A.G. (1987) *Manufacturing Engineering, Economics and Processes*, Prentice-Hall, Englewood Cliffs, NJ.

Schey, J.A. (1977) *Introduction to Manufacturing Processes*, McGraw-Hill, New York.

Suh, N.P. (1990) *The Principles of Design*, Oxford University Press, Oxford.

Waterman, N. and Ashby, M.F. (1997) *Chapman and Hall Materials Selector*, Chapman and Hall, London, Chapter 1.6.

Yankee, H.W. (1976) *Manufacturing Processes*, Prentice-Hall, Englewood Cliffs, NJ.

Cost modelling

Clark, J.P. and Field, F.R. III (1997) 'Techno-economic issues in materials selection', in *ASM Metals Handbook*, Vol. 20. American Society for Metals, Metals Park, Ohio.